Busy Simulations

ITS Upgrade WP10 PRR

Simon Voigt Nesbo <svn@hvl.no>

Summary of ALPIDE busy mechanisms

- ALPIDE chip asserts BUSY_ON when:
  - 3 MEBs are in use in triggered mode
  - 2 MEBs are in use in continuous mode
  - Frame FIFO reaches ALMOST_FULL1 watermark (48 events)
  - **Assertion of BUSY_ON does not imply data loss**

- Busy violation:
  - Occurs when all 3 MEBs are in use, the chip is busy and it receives a new trigger
  - This can also happen in continuous mode (theoretically)
  - **Data is lost**

- Flushed incomplete:
  - Occurs in continous mode when:
  - 2 MEBs are already in use, the chip is busy and it receives a new trigger
  - Readout of the oldest MEB slice is stopped and data flushed
  - **Data is lost**

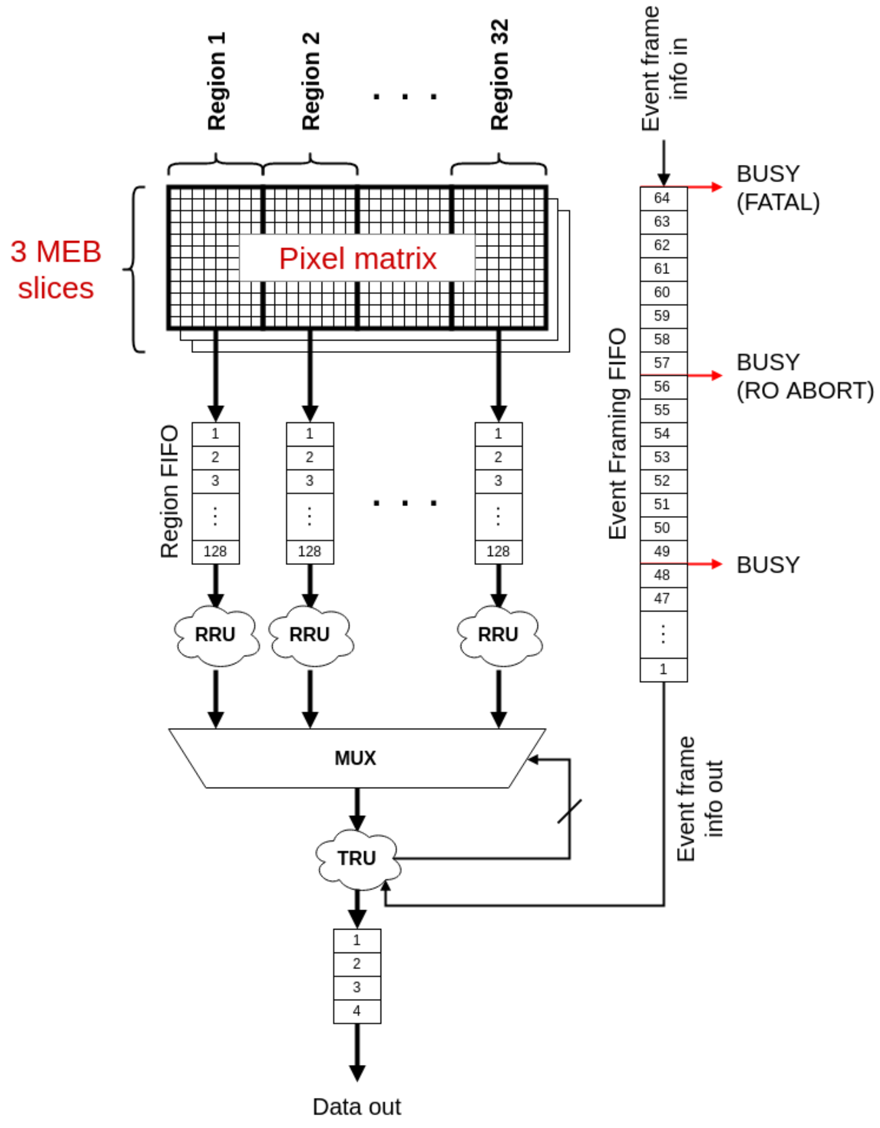

Summary of ALPIDE busy mechanisms

- READOUT_ABORT/DATA_OVERRUN mode:
  - Occurs when the frame FIFO reaches critical levels
  - Data discarded and empty frames transmitted
  - **Data is lost**

- FATAL mode:
  - Occurs when the frame FIFO overflows
  - All future data frames will be marked FATAL (until it is reset)
  - **Data is essentially garbage**

- The last two modes are not observed in these simulations at reasonable event rates

SystemC simulation setup

ALICE ITS UPGRADE

- Inner Barrel simulated
- Pixel dead time: 200 ns
- Pixel active time: 6000 ns

Parameter	PbPb triggered	PbPb continuous	pp continuous
Event rates	50, 100, 200 kHz	50, 100, 200 kHz	1, 5, 10 MHz
Number of interaction events	20,000	20,000	100,000
MC event pool	10k MB-PbPb + 100k QED	10k MB-PbPb + 100k QED	10k MB-pp
Strobe length	100 ns	4.9, 9.9, 19.9 us	4.9, 9.9, 19.9 us
Strobe gap	N/A	0.1 us	0.1 us
Trigger filter time	1230 ns	N/A	N/A
Minimum busy cycles	8 (200 ns)	160, 360, 760 (4, 9, 19 us)	160, 360, 760 (4, 9, 19 us)

PbPb triggered mode results

PbPb simulations (triggered mode) – Results

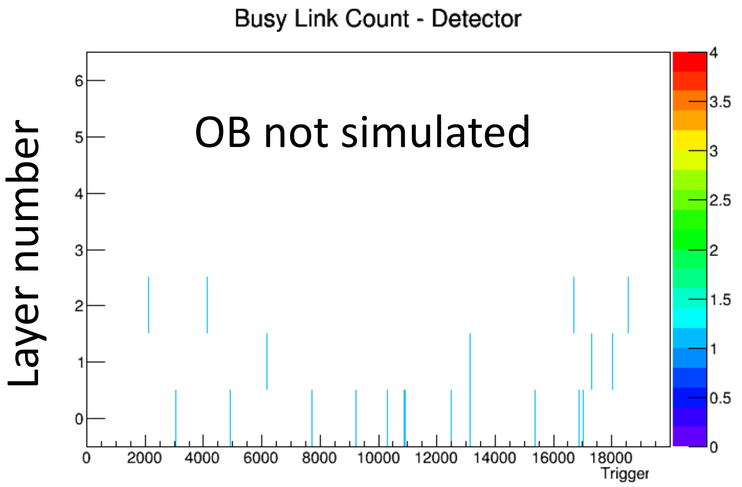

BUSY vs trigger, 50kHz triggered mode

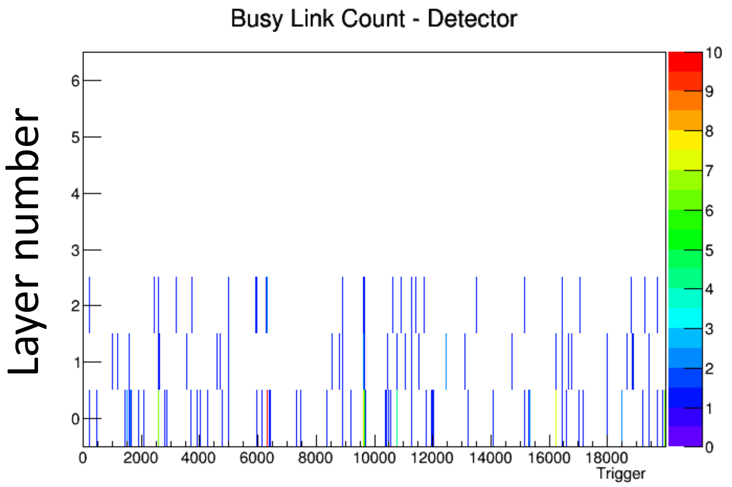

BUSY vs trigger, 100kHz triggered mode

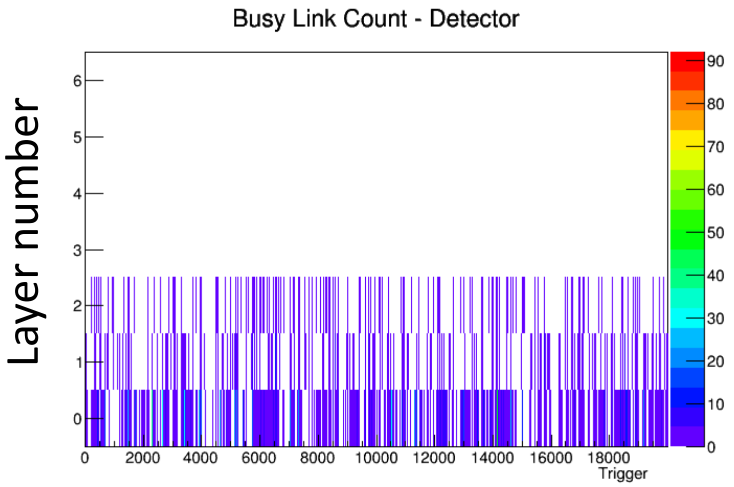

BUSY vs trigger, 200kHz triggered mode

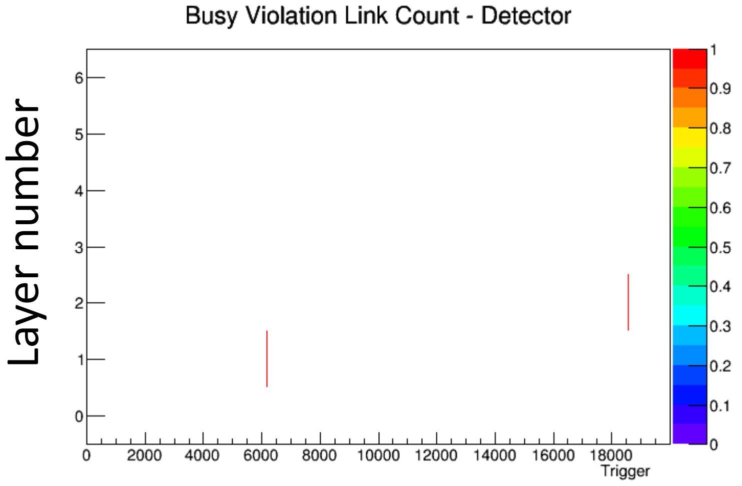

BUSYV vs trigger, 50kHz triggered mode

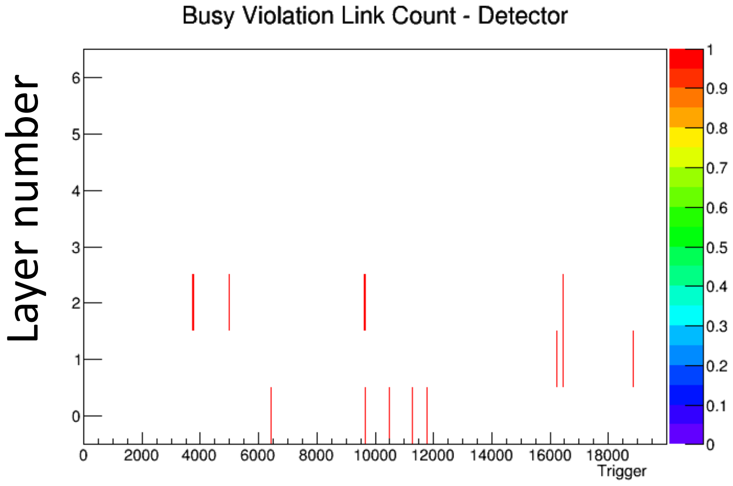

BUSYV vs trigger, 100kHz triggered mode

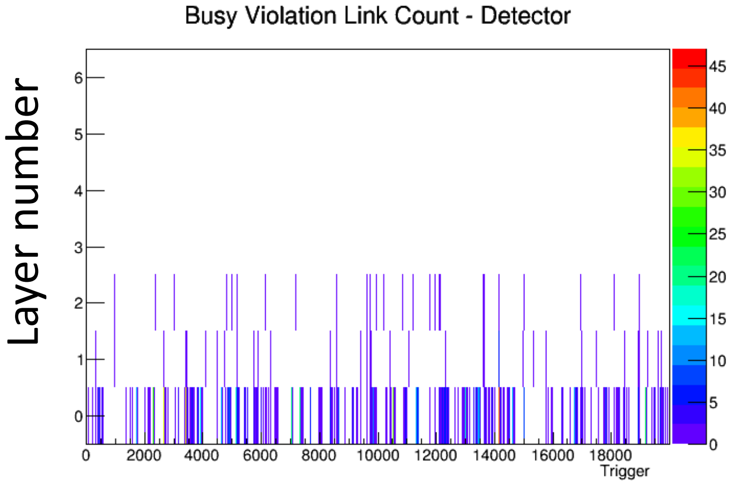

BUSYV vs trigger, 200kHz triggered mode

PbPb simulations (triggered mode) – Results

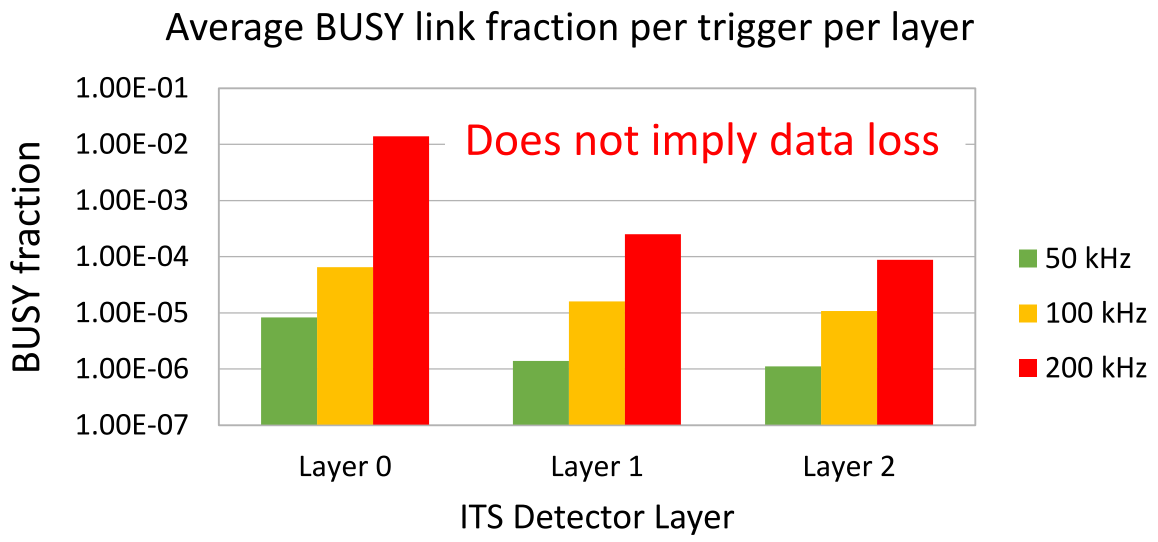

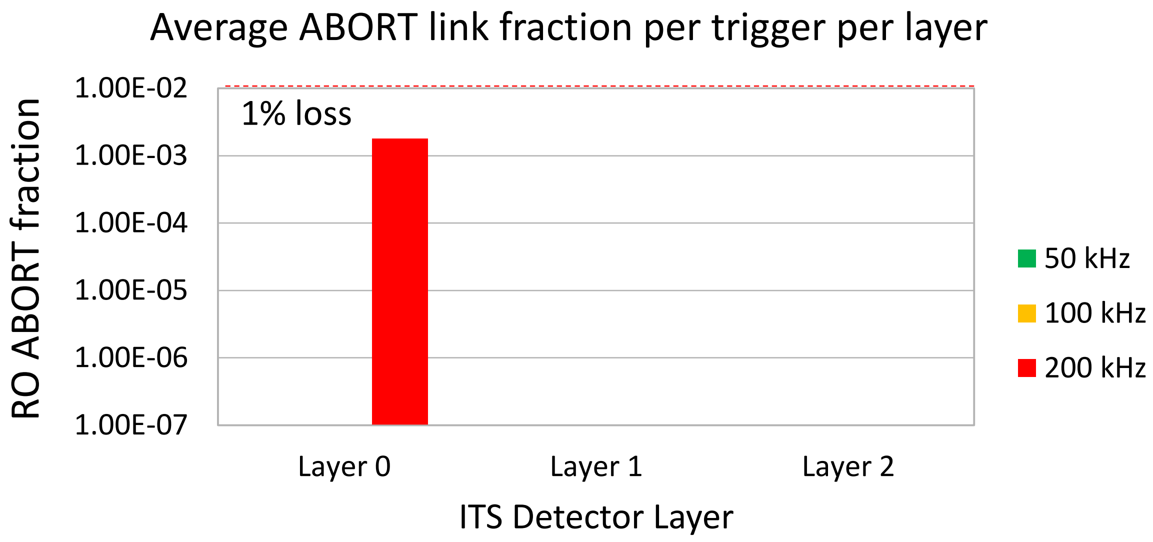

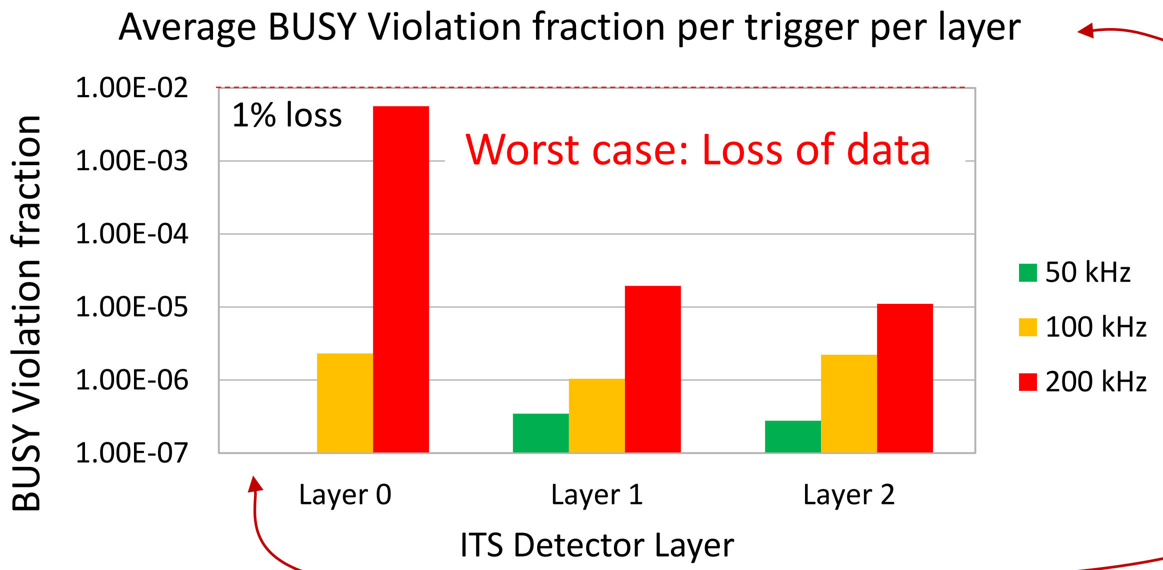

(1.0 – BUSYV fraction) is essentially the efficiency per layer (assuming that ALPIDE frames lost due to BUSYV is a relatively accurate representation of data loss)

Note:
The statistics simulated has a large error, because only 20,000 events were simulated. At 50 kHz there is 0 BUSYV in layer 0, and 1 BUSYV each in layer 1 and 2.

PbPb continuous mode results

PbPb simulations (continuous mode, 100kHz) – Results

ALICE ITS UPGRADE

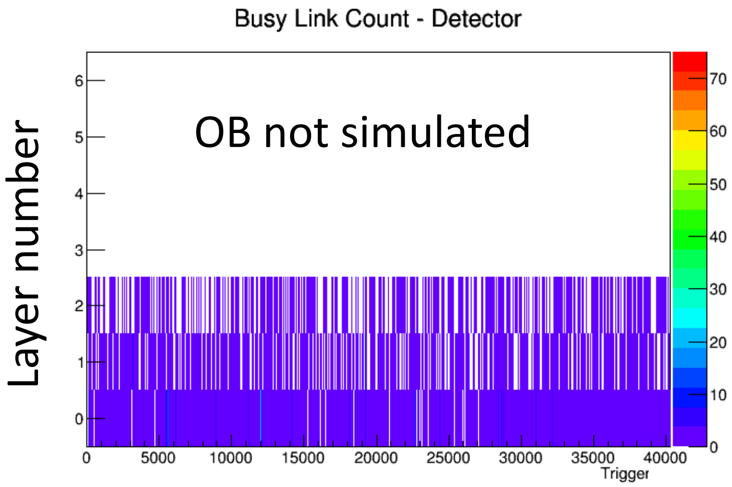

BUSY vs trigger, 5us strobe

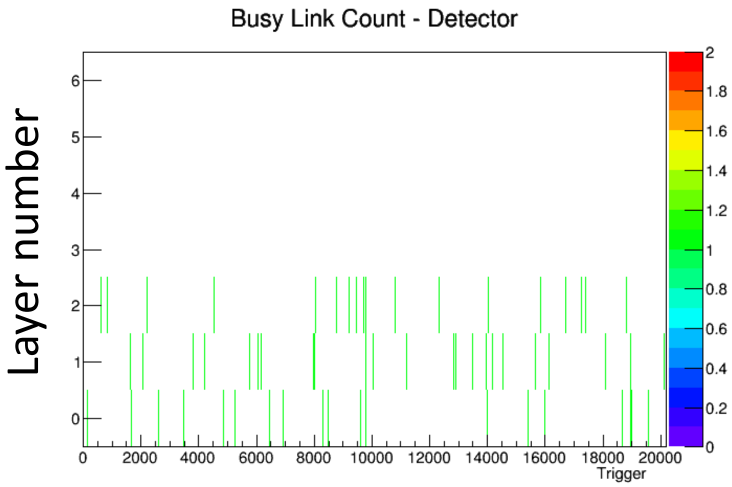

BUSY vs trigger, 10us strobe

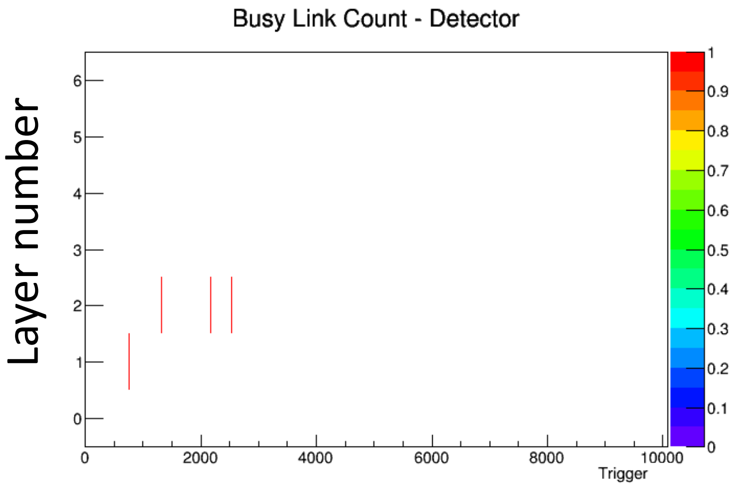

BUSY vs trigger, 20us strobe

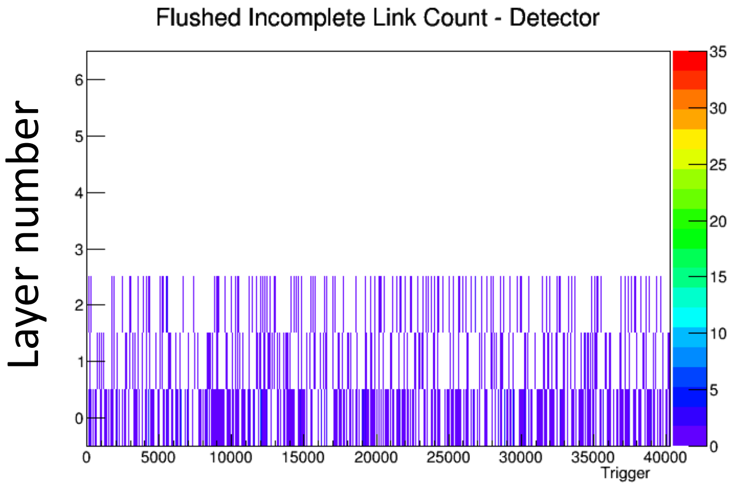

FLUSH INCOMPL. vs trigger, 5 us strobe

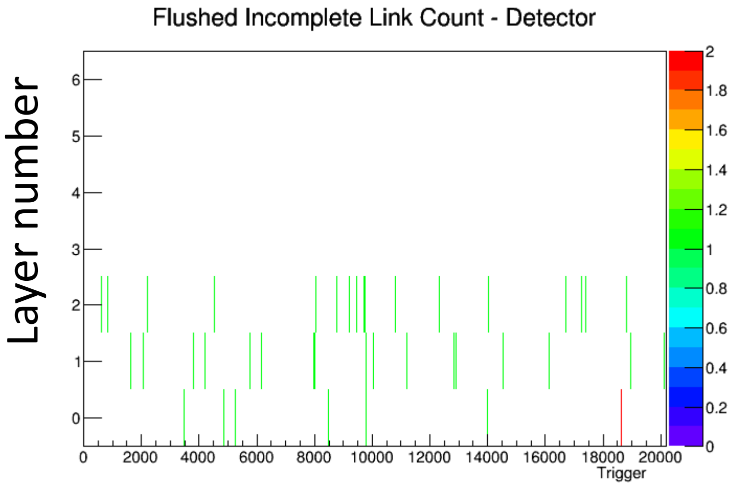

FLUSH INCOMPL. vs trigger, 10us strobe

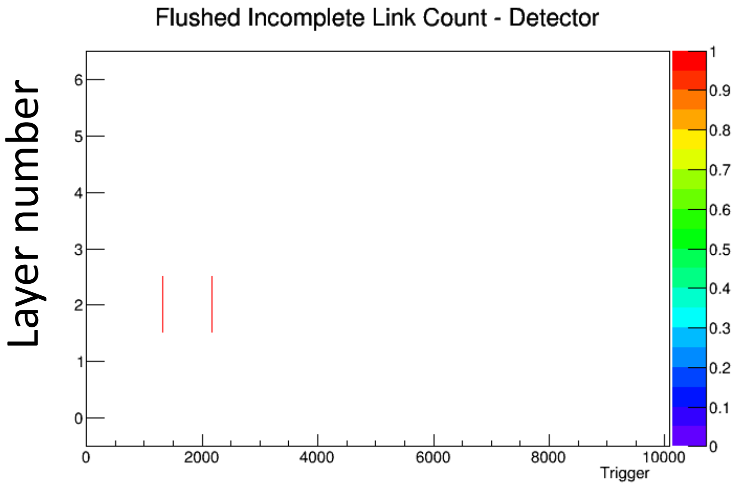

FLUSH INCOMPL. vs trigger, 20us strobe

PbPb simulations (continuous mode) – Results

Average BUSY link fraction per trigger per layer

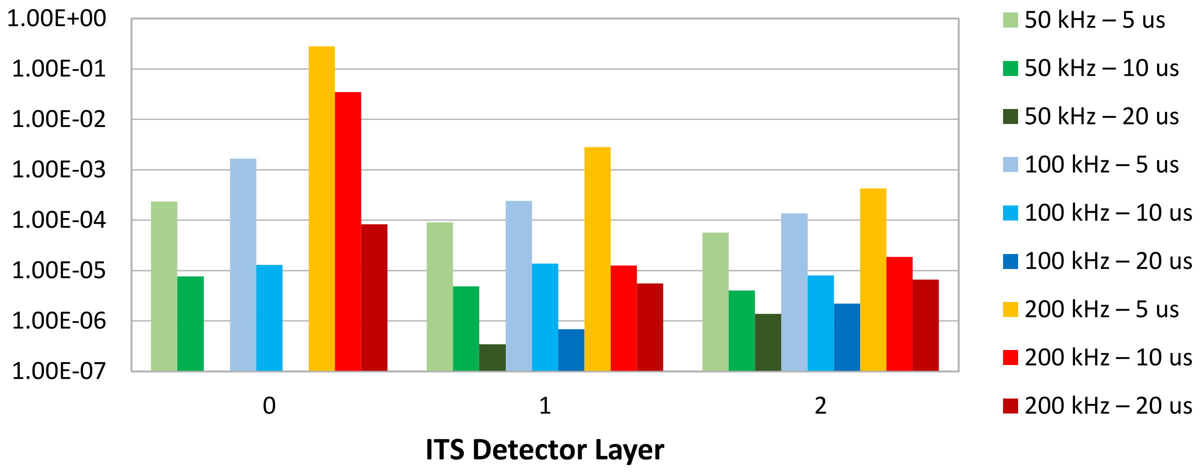

Average ABORT link fraction per trigger layer

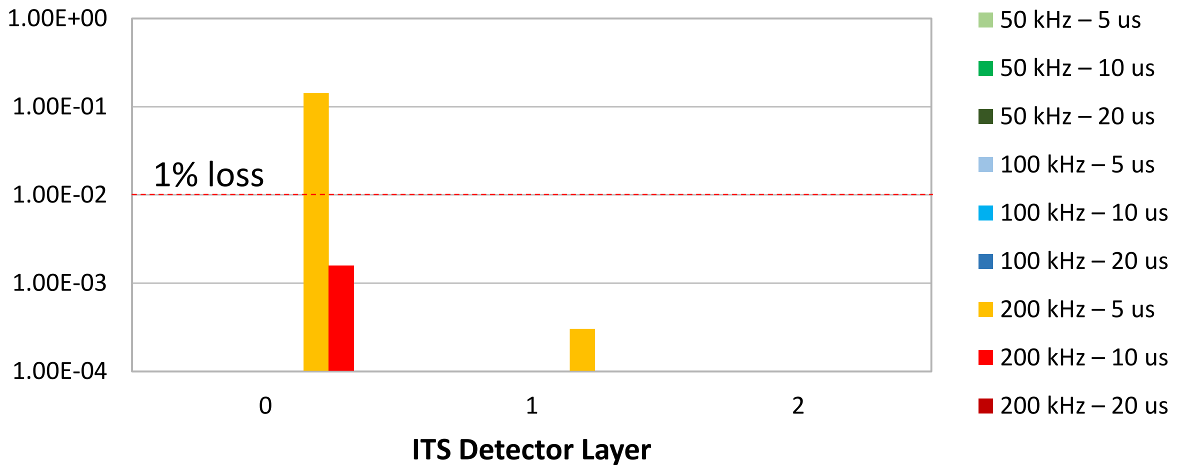

Average BUSYV link fraction per trigger per layer

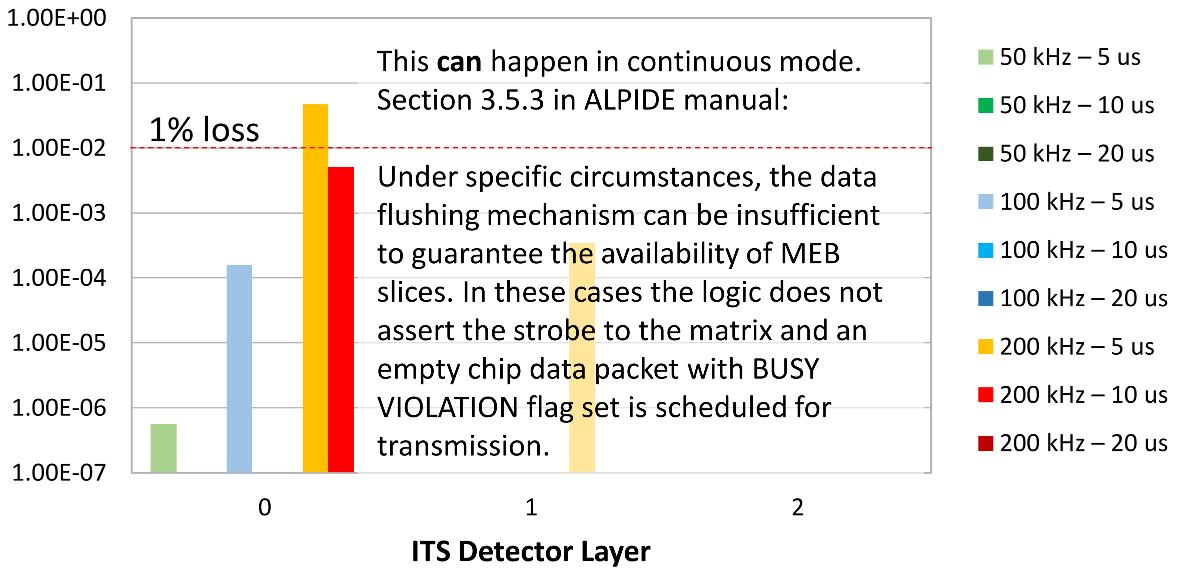

Average FLUSH link fraction per trigger per layer

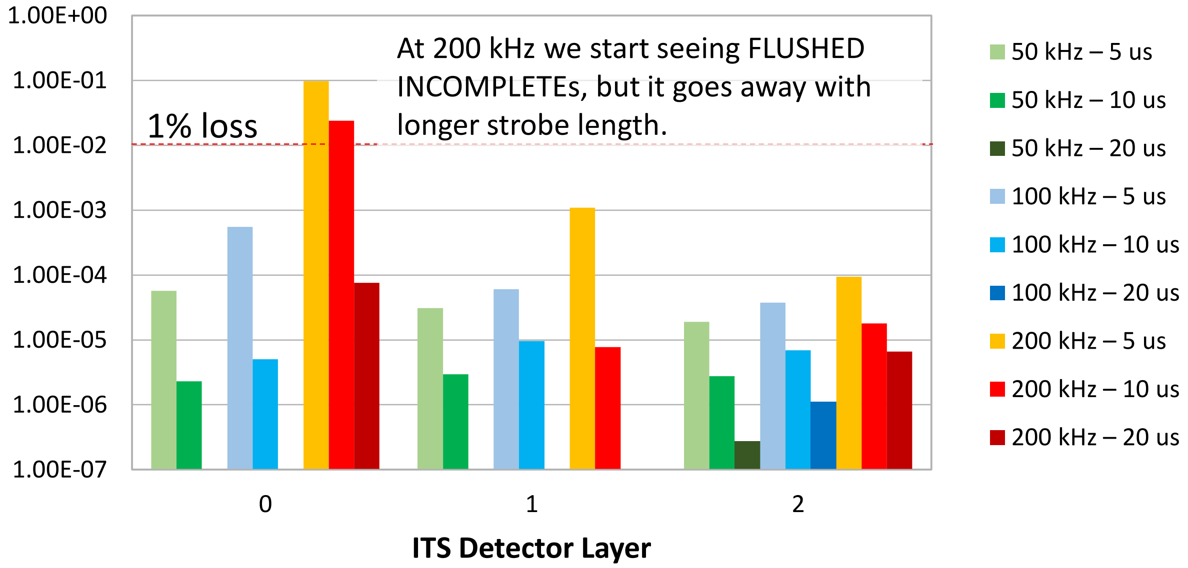

pp continuous mode results

pp simulations (continuous mode, 5 us strobe) – Results

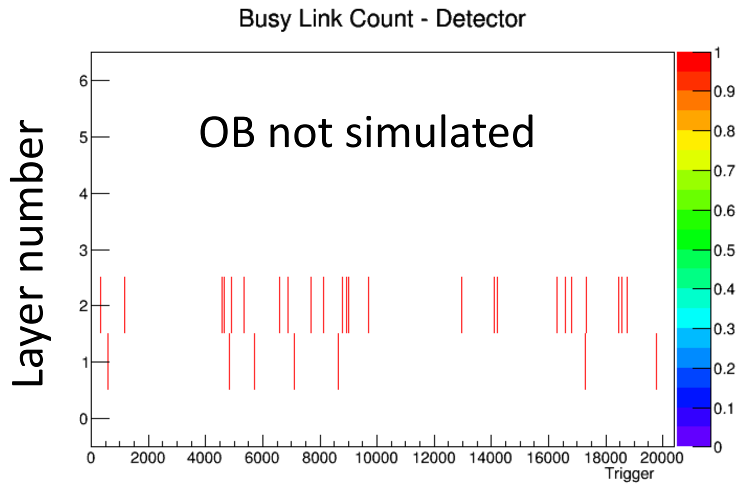

BUSY vs trigger, 1 MHz

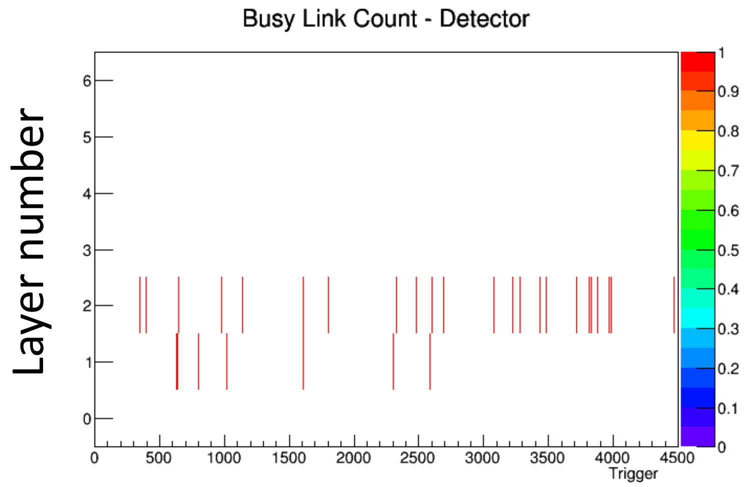

BUSY vs trigger, 5 MHz

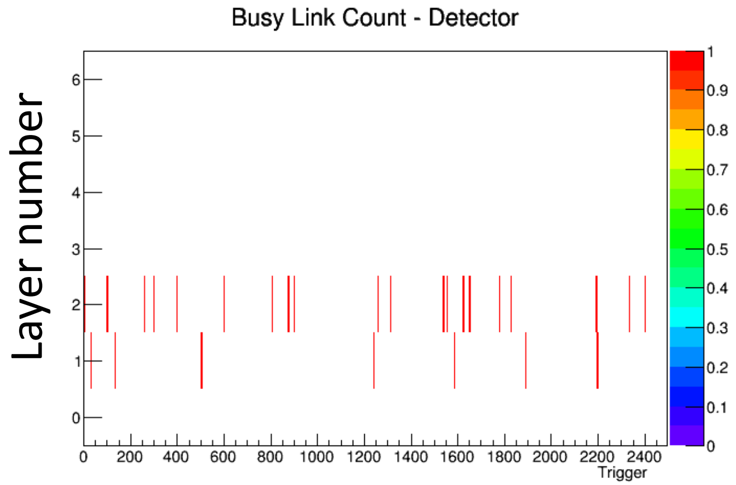

BUSY vs trigger, 10 MHz

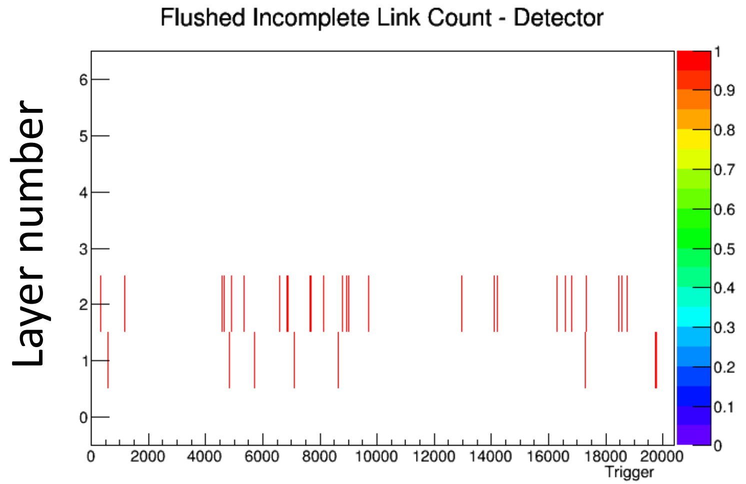

FLUSH INCOMPL. vs trigger, 1 MHz

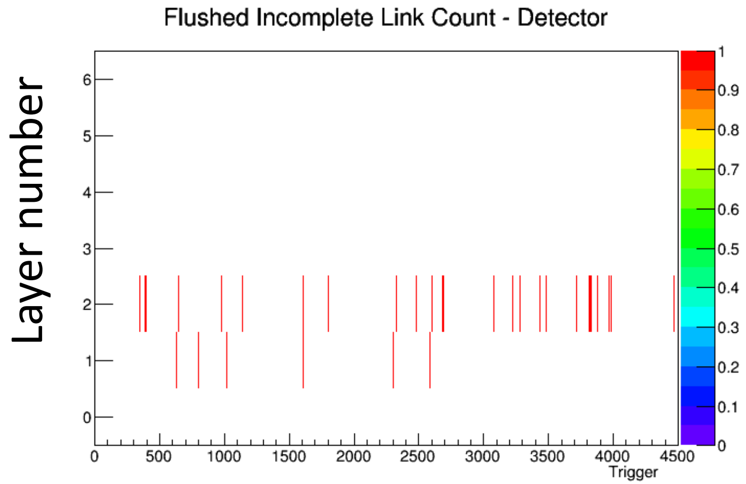

FLUSH INCOMPL. vs trigger, 5 MHz

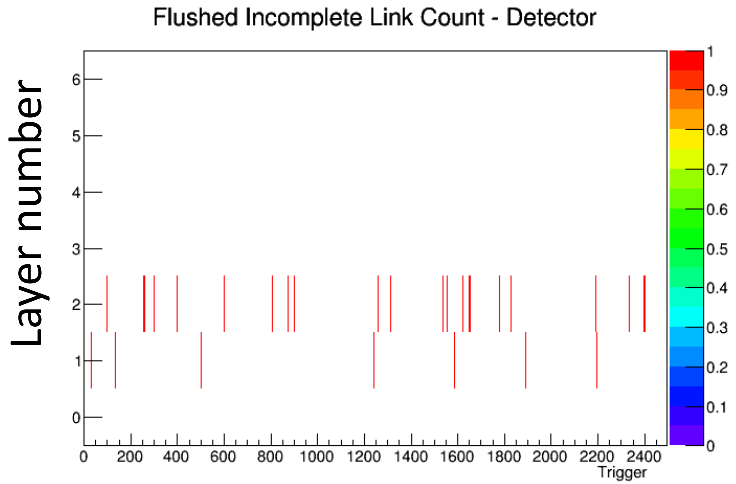

FLUSH INCOMPL. vs trigger, 10 MHz

No BUSY or FLUSHED INCOMPLETE seen at 10 us and 20 us strobe lengths

pp simulations (continuous) – Results

ALICE ITS UPGRADE

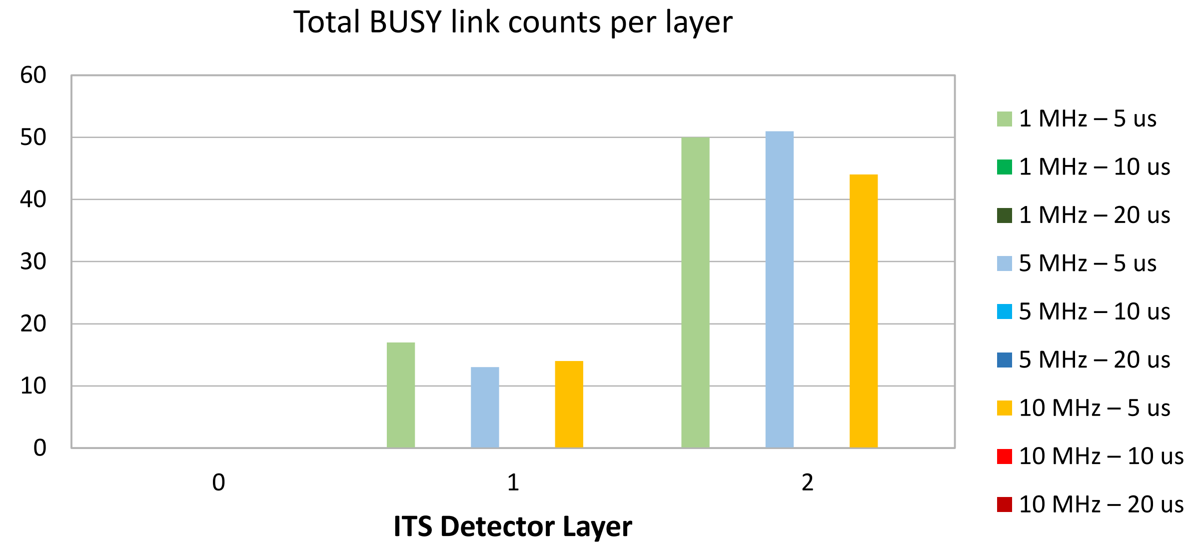

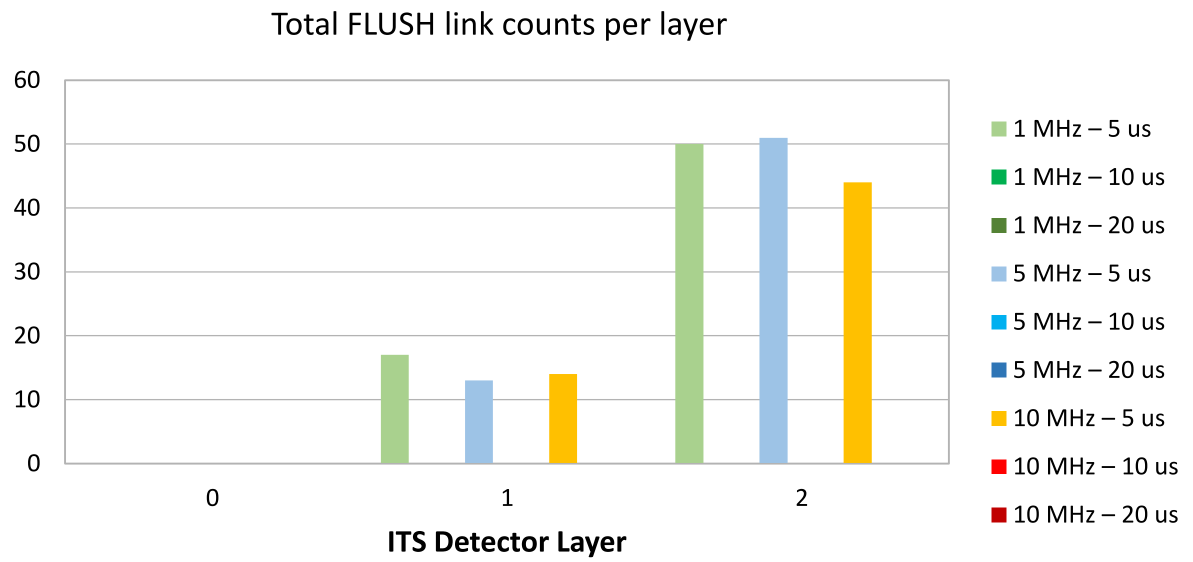

- The number of BUSY/COUNT observed is relatively constant for a given number of events
  - Independent of event rate or strobe length
- This is because of **3 atypical interaction events** in the pp MC event pool with high multiplicity for specific chips, that due to limited statistics have a very high impact on these simulations and plots
  - They always lead to busy and flushed incomplete
- This has also been observed in other simulations at a number of different event rates
- **The problem goes away with longer strobe (10 us and 20 us)**
- **The problem is not due to occupancy**

Conclusions

- Busy is not critical for nominal event rates
  - ALICE: 50 kHz PbPb, 200 kHz pp
  - ITS Upgrade: 100 kHz PbPb, 400 kHz pp
- A flush/busyv in one link now and then has little impact on the overall readout efficiency of the detector
  - Busy violation/flush incomplete count below acceptance level at nominal rates
  - **Monitoring busy/busyv/flush/etc seems sufficient at nominal event rates**
- The **only action** a possible busy system can take is to withhold the next trigger
- With a busy system, ITS could probably cope with higher rates in triggered or continuous with short strobe (5 us) and still extract quality data
- Alternatively, higher rates can be run in continuous mode with long strobe (e.g. 20 us), **at the cost of more pile-up and reconstruction challenges**

Possible Busy System

Busy timing diagrams

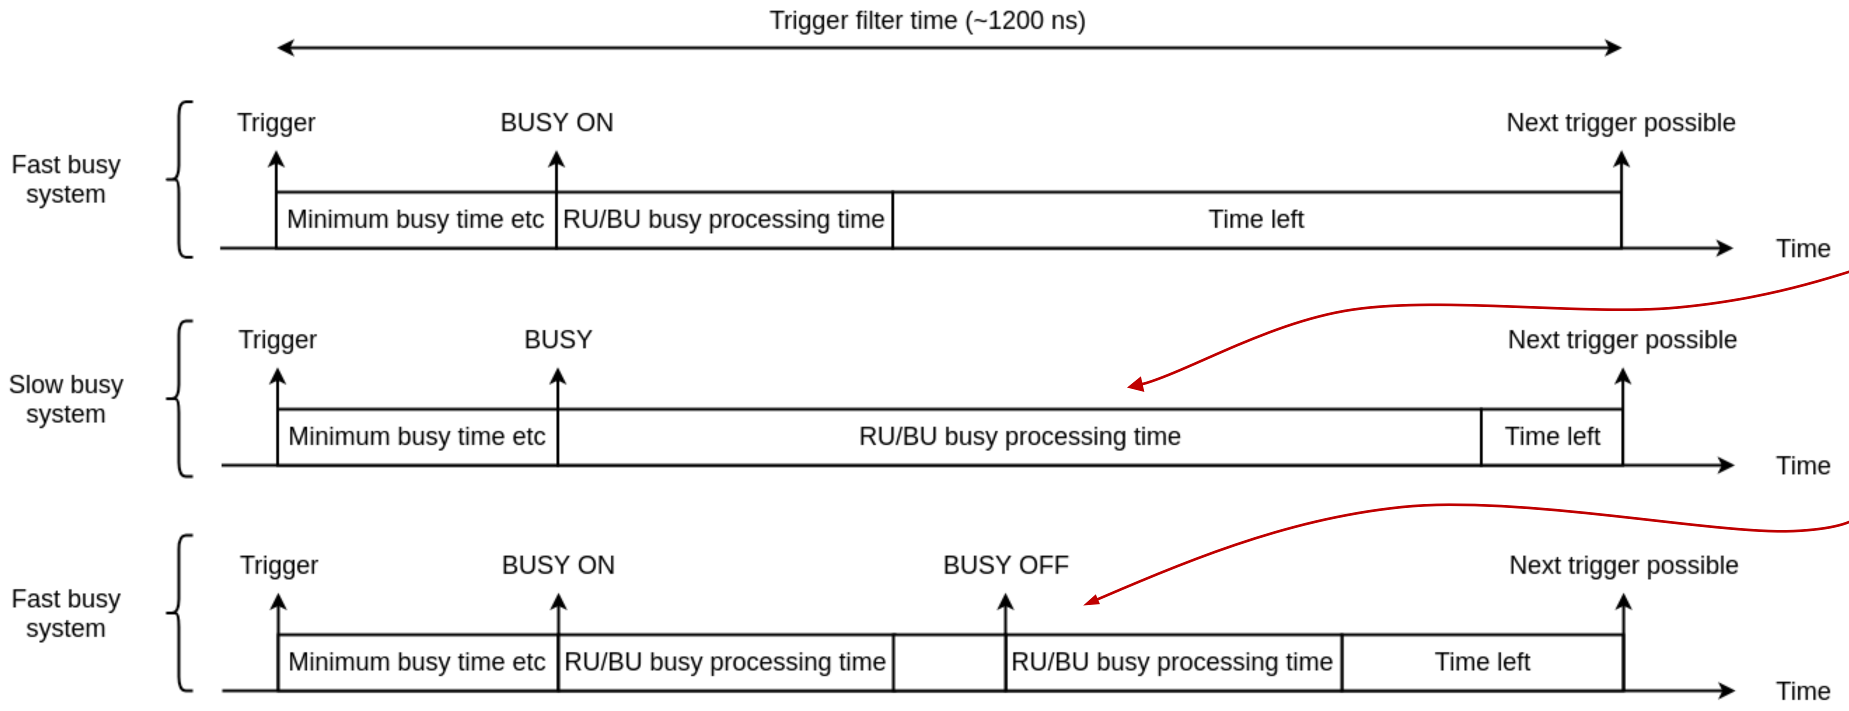

If processing time is very long, there is no time to make a new decision if BUSY_OFF comes in

With short processing time, we are able to make a new decision on BUSY_OFF

- Busy subsystem must be faster than trigger filtering time to be of any use
- From simulations we know that a fraction of busies lead to busy violations/flush
- Hence, **most BUSY_ON will be followed by BUSY_OFF before next trigger**
- Busy subsystem needs to consider the BUSY_OFF
- **Busy subsystem should be as fast as possible, so it can act on the most recent information (BUSY_OFF)**

Busy System Conceptual Idea

ALICE ITS UPGRADE

- Two concepts for a busy system have been considered:
  - Daisy chained Reaodut Units (no dedicated Busy Unit)
    - Slower, with synchronization issues
  - Dedicated «Busy Unit»
    - Star topology
    - No synchronization issues
- There is room in the mini crates for Busy Units
- Firmware implementation:
  - Count number of busy ALPIDE data links in each RU
  - Global busy status/action based on threshold of busy links
  - Hold back trigger on global busy status

Busy Unit based on modified Readout Unit design

- Use a new transition board to accept direct busy input from several RUs
- Existing RU has only receiver lines to transition board
- BU design must have transmitter lines routed to transition board

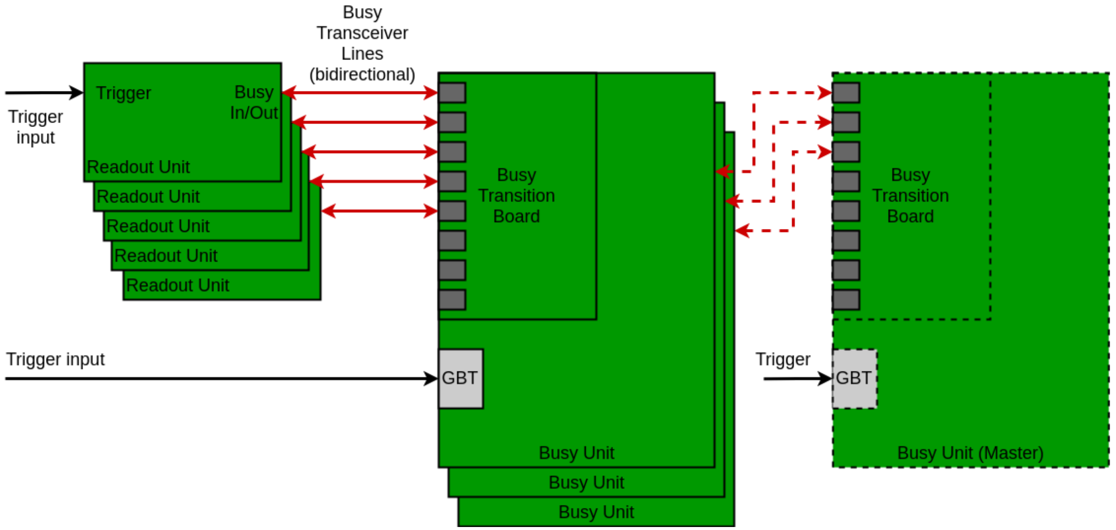

Backup

SystemC model

- Simulation model of readout chain in Alpide chip
- Initially intended for busy simulations only
- More general purpose at this point

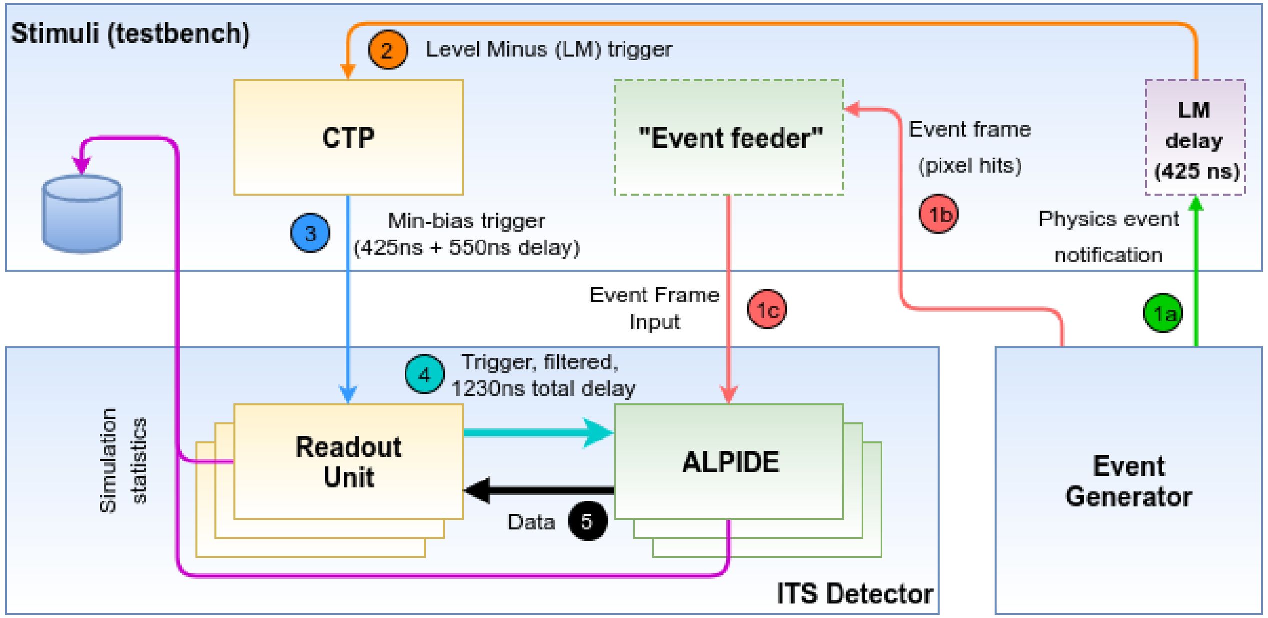

Accuracy of SystemC model

- Aims to be a relatively accurate model, close to a 1:1 copy of real chip in SystemC
- Top Readout Unit (TRU) and Region Readout Unit (RRU)
  - Full model of FSMs, based on diagrams from Alpide EDR presentation
- Custom clustering method in C++/SystemC, based on interpretation of Alpide manual and data format
- Frame ReadOut and Management Unit (FROMU)
  - No direct counterpart to FROMU in code, but similar functionality implemented
- Data Management Unit (DMU)
  - Currently no direct counterpart to DMU
- There is a 4-word deep FIFO representing DMU FIFO
- Data Transfer Unit (DTU)
  - No DTU in the model. Serializing and encoding not necessary for our purposes (busy simulations)
  - DTU «implemented» with a dummy delay

PbPb simulations – Event generation

A pool of discrete MC events for PbPb and QED were generated using the itsuTestBench in AliRoot:

- Using Hijing for random particle generation, and GEANT4 for tracking and detector response
  - Random events were picked from the pool using a uniform distribution, and fed to the detector in SystemC
  - Random time between events is picked from an exponential distribution, with lambda = 1 / avg_event_rate (time)
- QED events generated with 250 ns integration time
  - Luminosity: 6E27, Integration time: 250 ns
  - QED event as input to the chips continuously at 250 ns intervals, independent of triggers and interaction events

PbPb simulations – QED hits

ALICE ITS UPGRADE

Triggered mode

Pulse shaping time (before threshold) and trigger delay left out for simplicity

Bundling 5 us of QED with each PbPb event leads to overestimation of QED background?

Separating them underestimates it at higher event rates (QED depends on luminosity?)

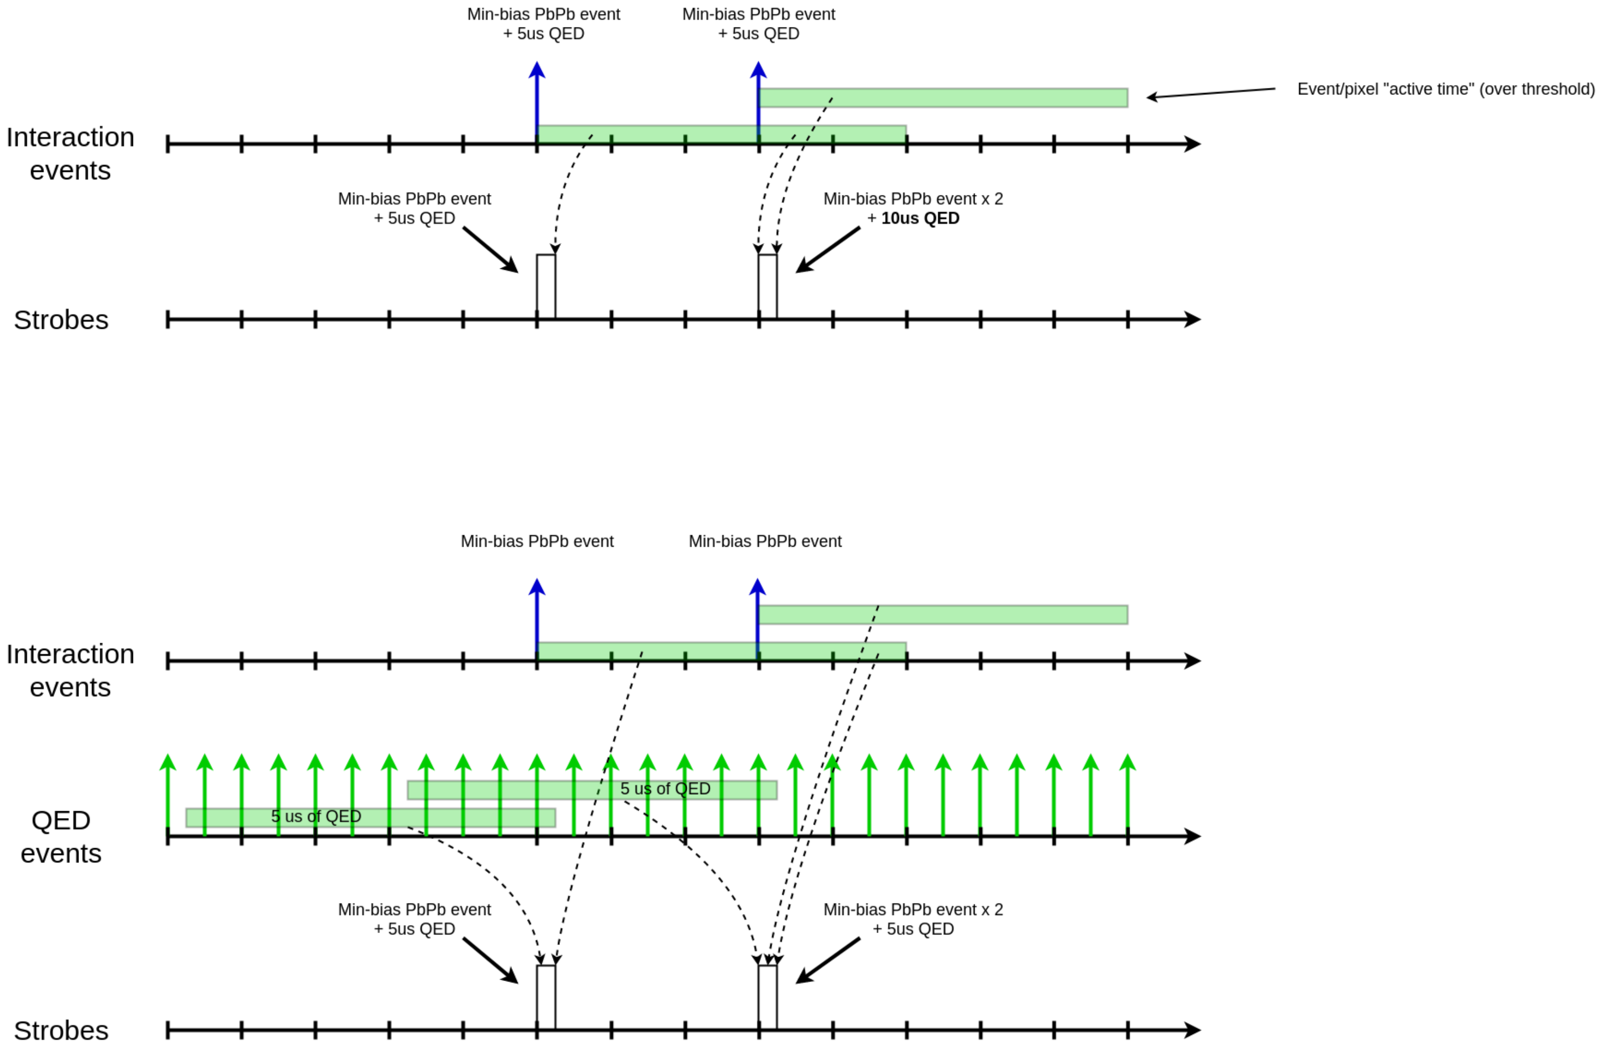

Busy stats recorded in simulation model

Scenario #1:
A link goes in and out of busy between triggers

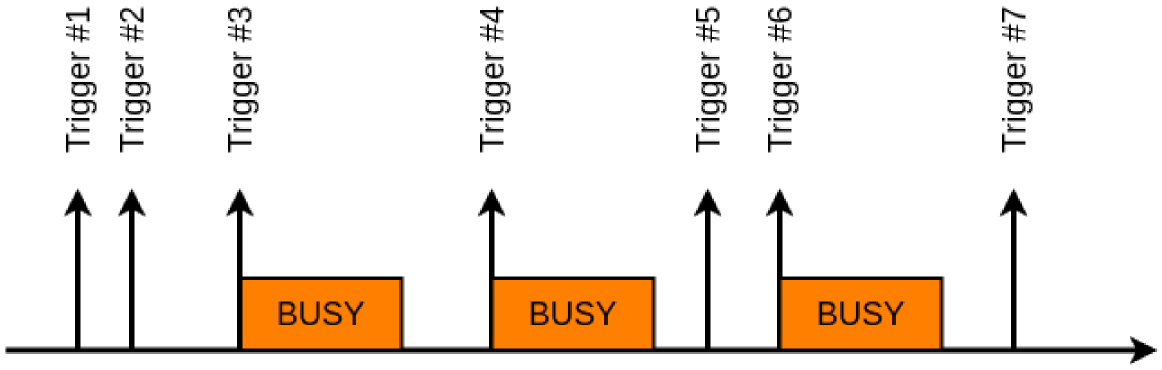

3 BUSY
0 BUSYV

Recorded by simulation

Scenario #2:
A link goes busy, and stays busy for more than one trigger, leading to busy violations

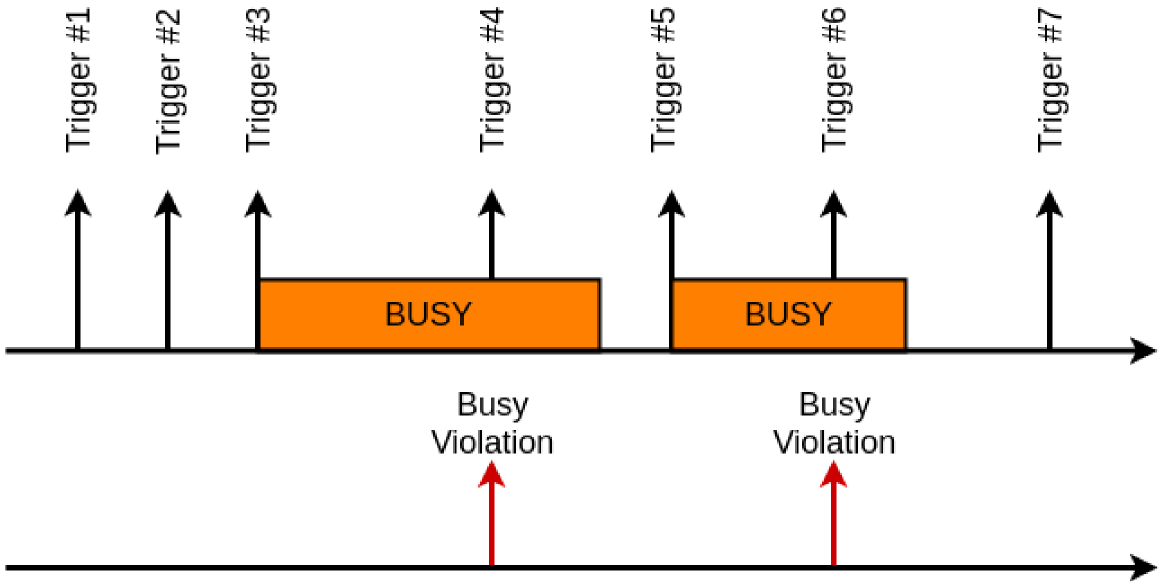

4 BUSY
2 BUSYV

Busy stats recorded in simulation model

ALICE ITS UPGRADE

- Total counts of triggers with BUSY, BUSYV, FLUSH, ABORT, FATAL is stored by simulation, summed for each ALPIDE data link, for each layer:

Num_triggers	40285				
Layer	BUSY	BUSYV	FLUSH	ABORT	FATAL
0	7199	687	2405	0	0
1	1393	0	349	0	0
2	980	0	273	0	0

Example data set:
100 kHz PbPb, 5 us strobe
Continuous mode

- Fraction of busy ALPIDE data links in layer:
  (BUSY count)/(Num_triggers*num_chips_in_layer)
- Same calculation for the other variables
- Acceptance = BUSYV fraction in layer

Minimum busy time

- There is a setting for “minimum busy width” in the ALPIDE chip:
  - Register **0x001B - Minimum Busy Width**
  - Number of clock cycles the chip has to be in busy state internally before it asserts BUSY_ON
  - Defaults to 8 clock cycles (ie. 200 ns)

- In triggered mode, the time the busy subsystem has to make a decision is:
  - trigger_filtering_time – minimum_busy_time = 1230 ns – 200 ns = 1030 ns

- In continuous mode, with 5 us strobe, the time is:
  - 5000 ns – 200 ns = 4800 ns

- If the busy subsystem is designed to make an action in 1 us, it makes sense to increase minimum busy time in continuous mode. With 160 clock cycles (4000 ns):
  - 5000 ns – 4000 ns = 1000 ns

- With 10 us and 20 us strobe, 360 and 760 min busy cycles gives us the the busy signal 1 us before the next strobe

Busy timing diagram (detailed)

ALICE ITS UPGRADE

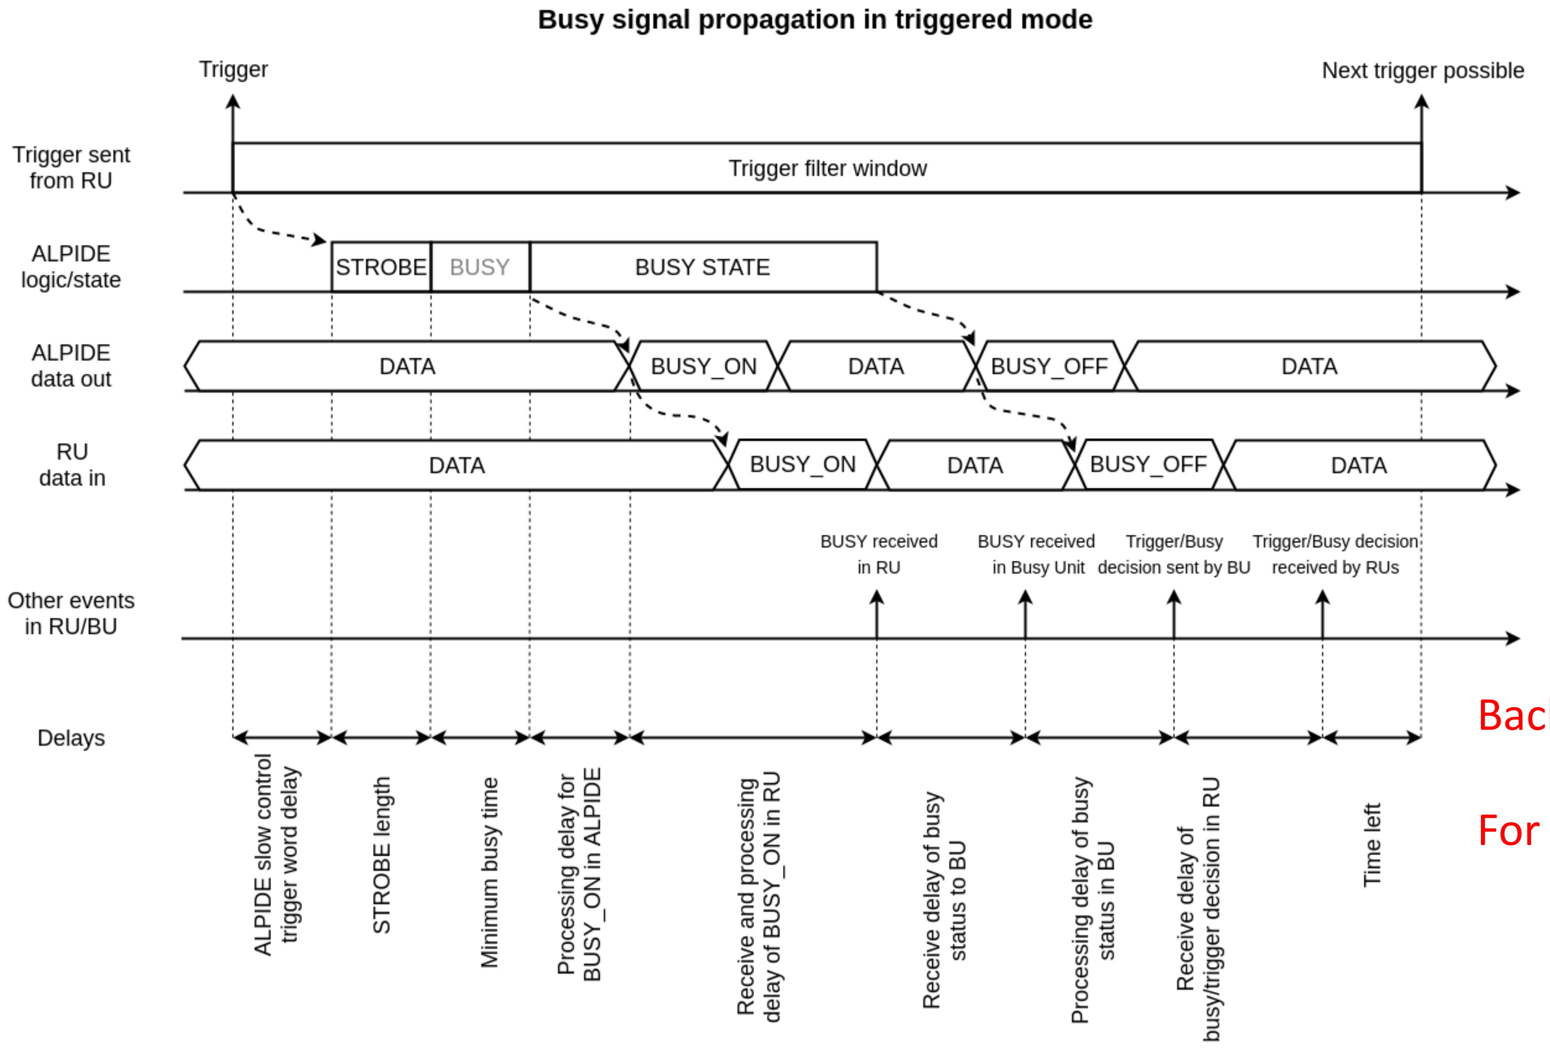

Backup

For mye detaljer

Effect of «missing chips» on reconstruction

- The level of systematic error on the acceptance is on the percentage-level[1]
- If the fraction of missing chips for an event is lower than 1E-3, it will not affect physics
- Since “missing events” will be marked with busy violation flag, it will be possible for reconstruction to take this into account
- Full reconstruction should start working with contributions from a minimum of 4 layers

[1] Private communication with Ruben Shahoyan

PbPb simulations (continuous mode, 50kHz) – Results

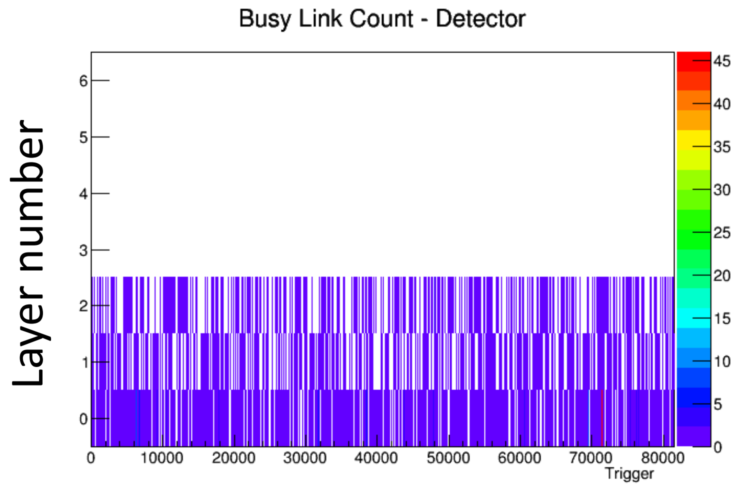

BUSY vs trigger, 5us strobe

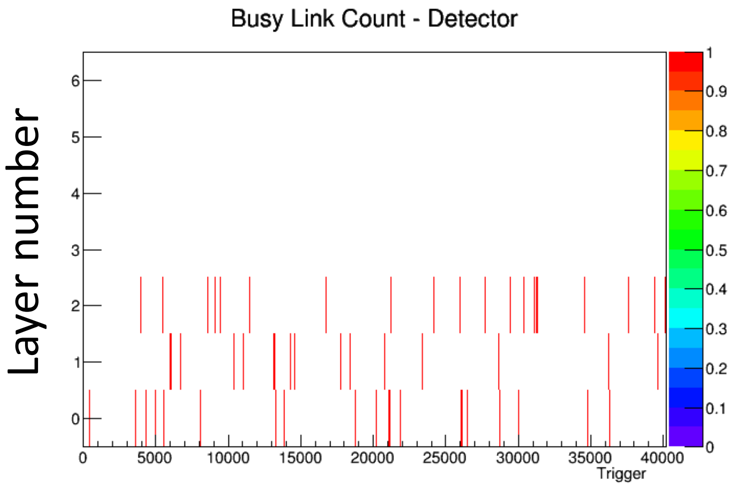

BUSY vs trigger, 10us strobe

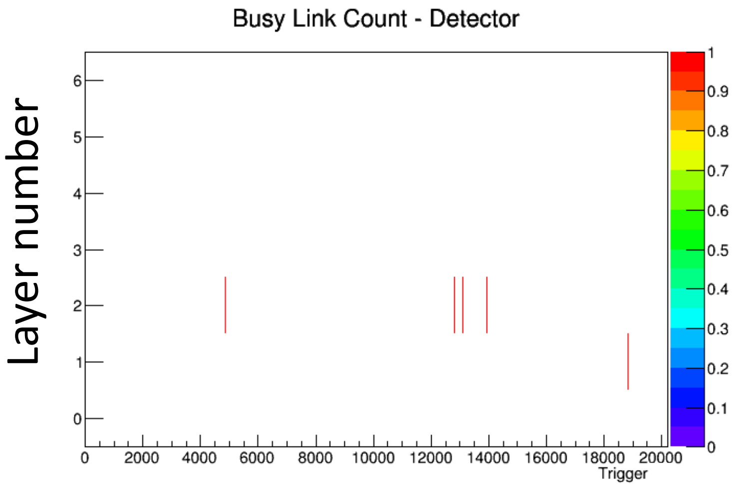

BUSY vs trigger, 20us strobe

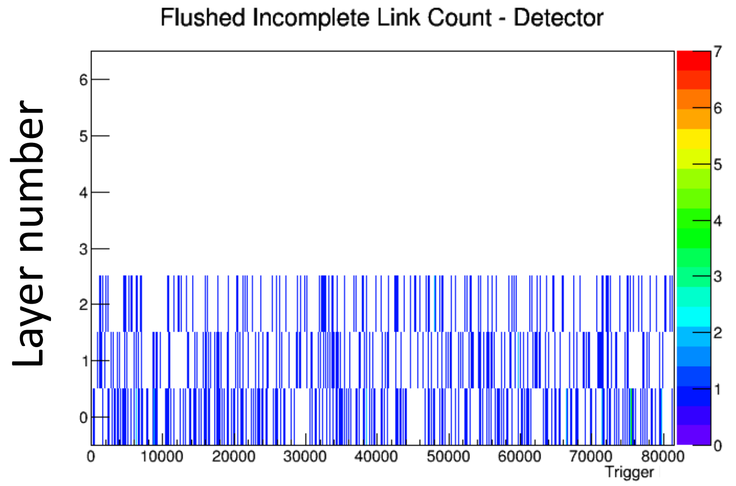

FLUSH INCOMPL. vs trigger, 5 us strobe

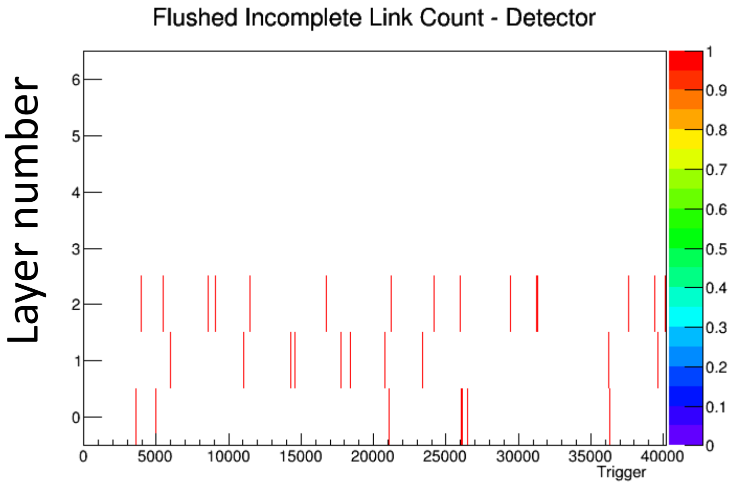

FLUSH INCOMPL. vs trigger, 10us strobe

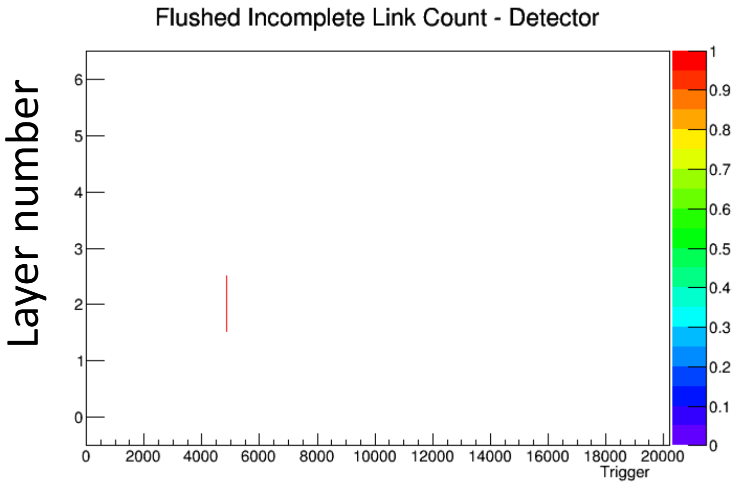

FLUSH INCOMPL. vs trigger, 20us strobe

PbPb simulations (continuous mode, 50kHz) – Results

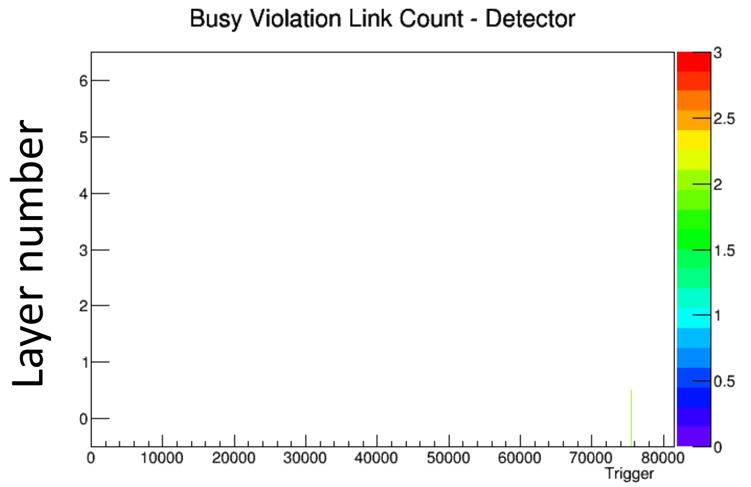

BUSYV vs trigger, 5us strobe

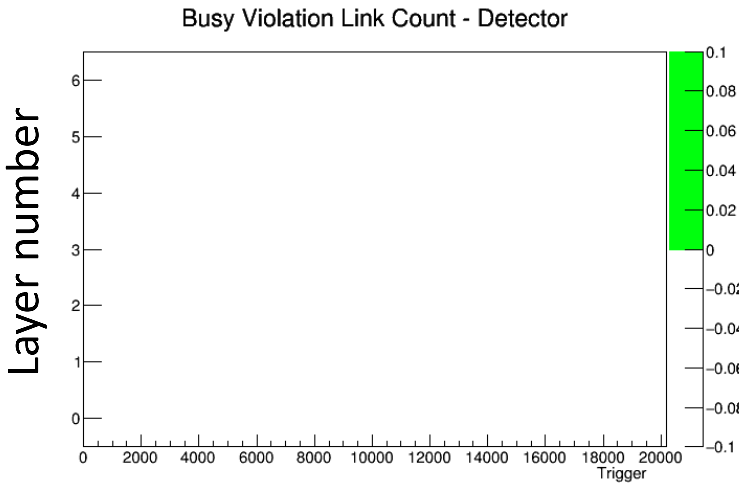

BUSYV vs trigger, 10us strobe

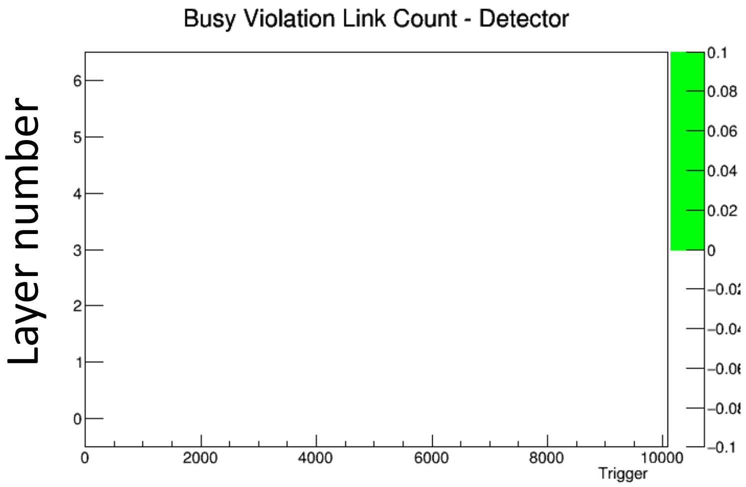

BUSYV vs trigger, 20us strobe

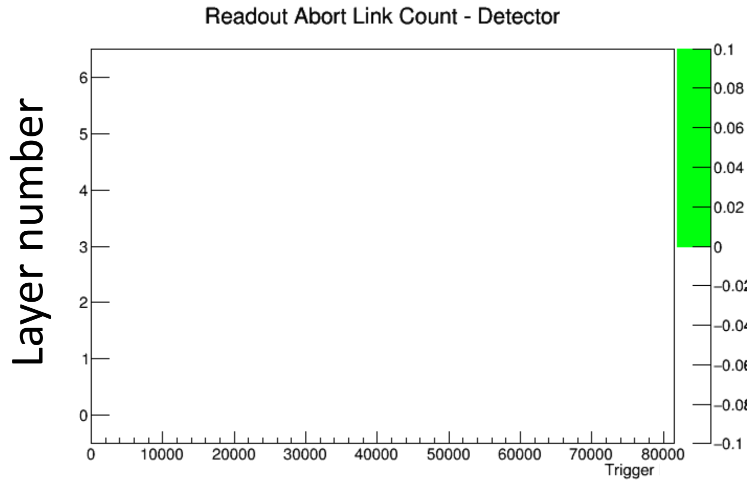

RO ABORT vs trigger, 5 us strobe

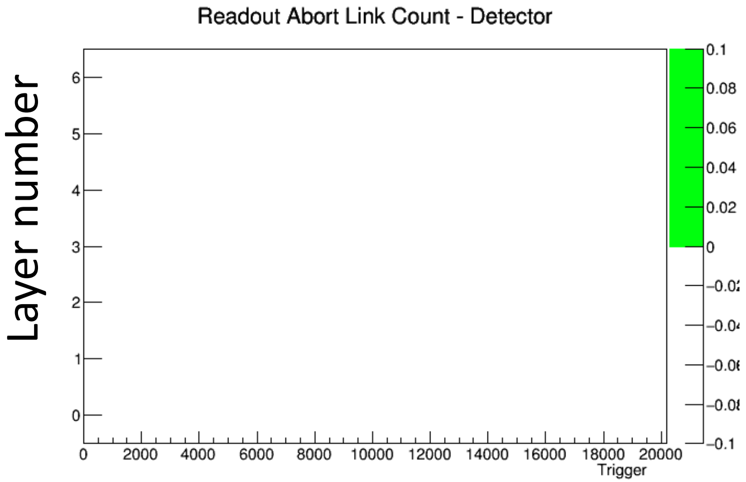

RO ABORT vs trigger, 10us strobe

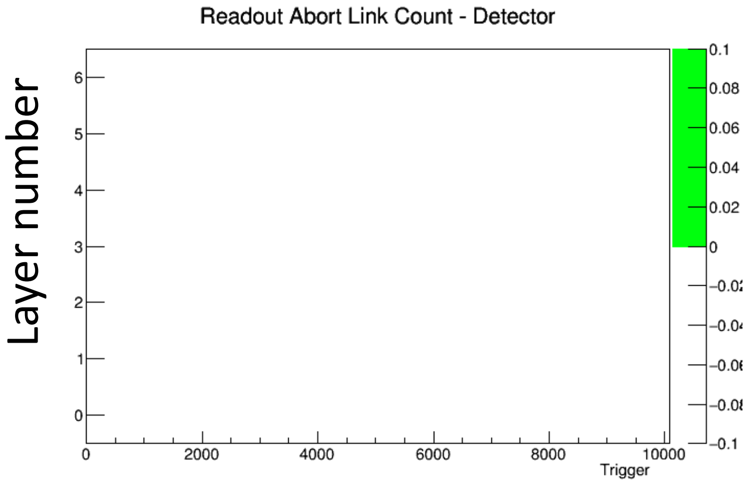

RO ABORT. vs trigger, 20us strobe

PbPb simulations (continuous mode, 100kHz) – Results

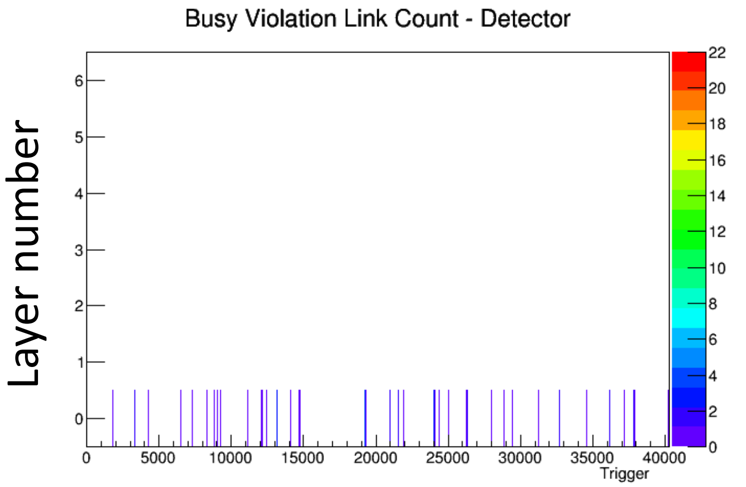

BUSYV vs trigger, 5us strobe

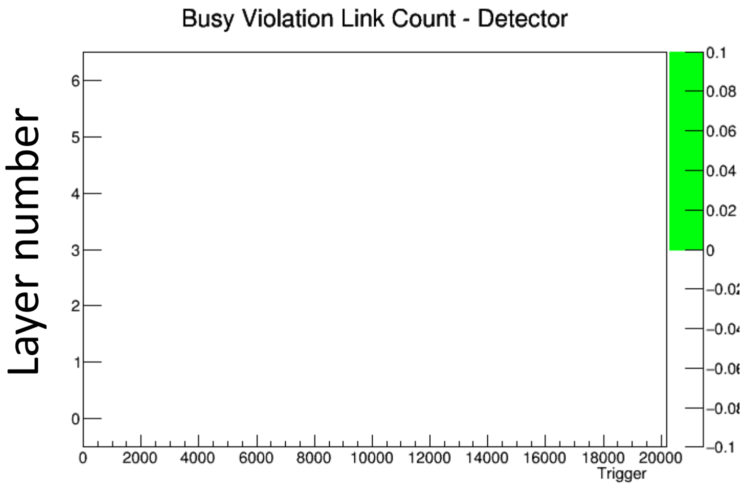

BUSYV vs trigger, 10us strobe

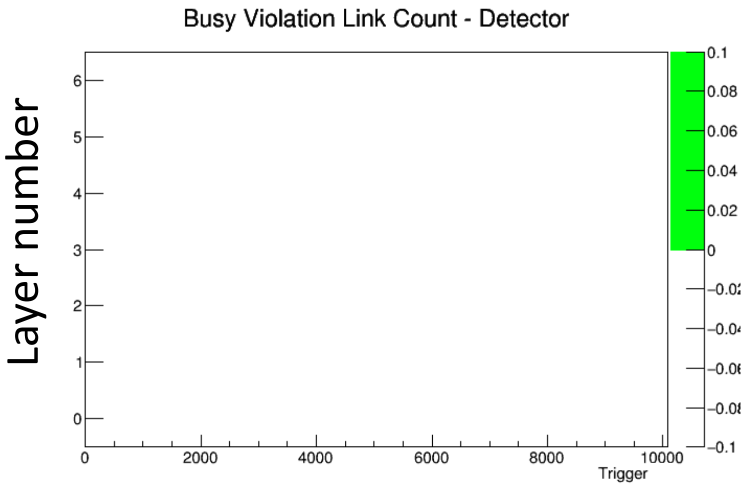

BUSYV vs trigger, 20us strobe

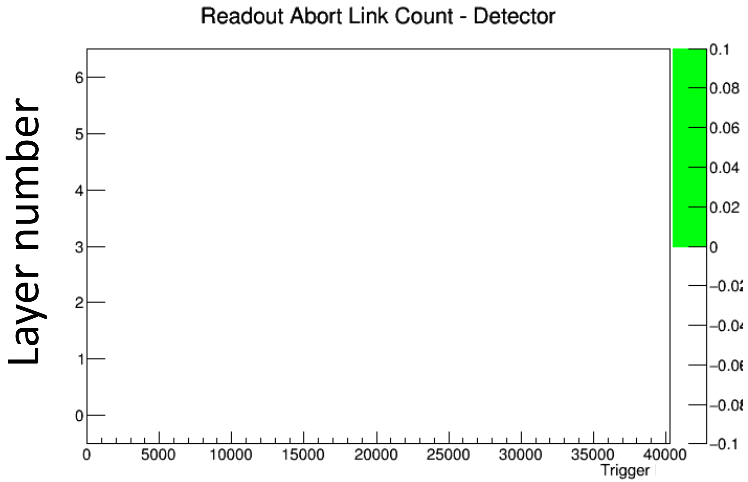

RO ABORT vs trigger, 5 us strobe

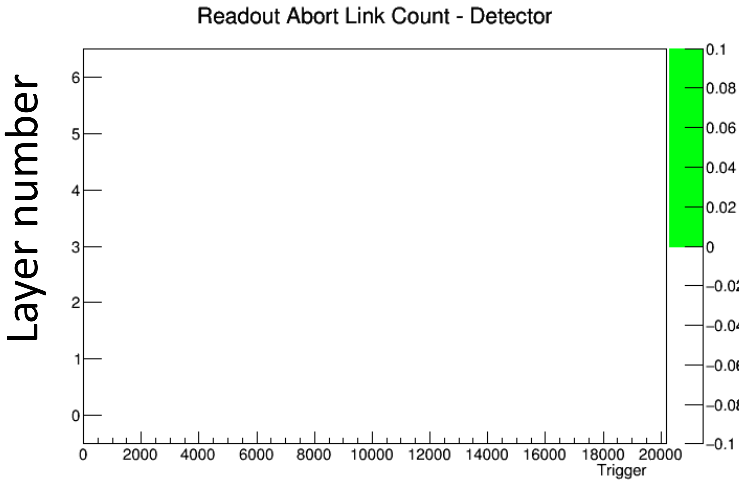

RO ABORT vs trigger, 10us strobe

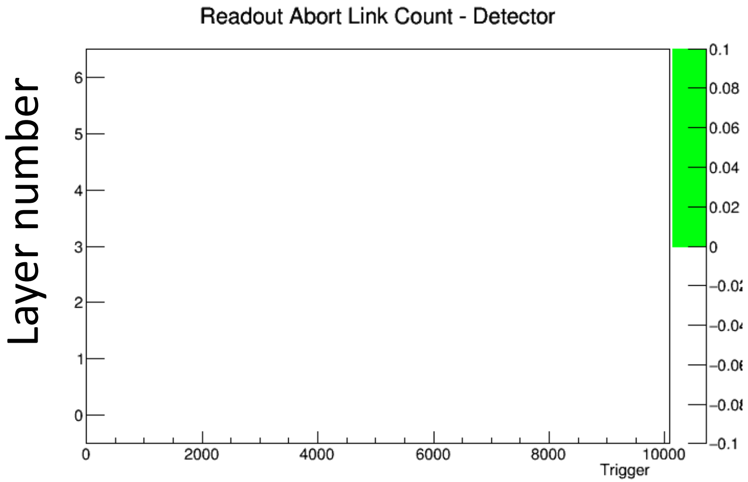

RO ABORT. vs trigger, 20us strobe

PbPb simulations (continuous mode, 200kHz) – Results

ALICE ITS UPGRADE

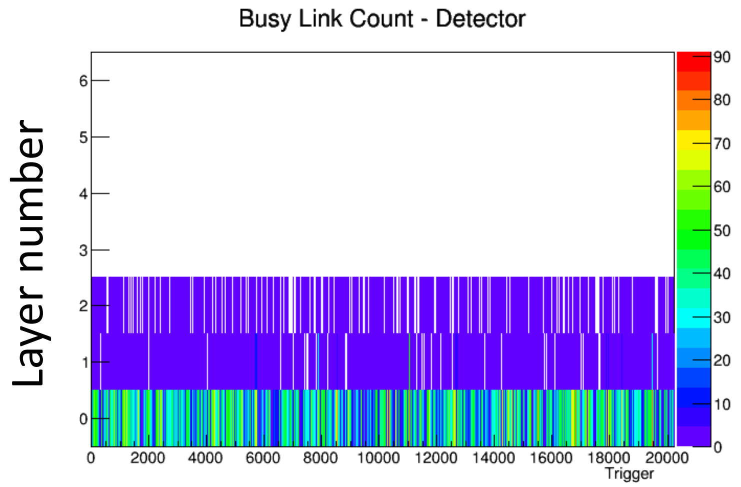

BUSY vs trigger, 5us strobe

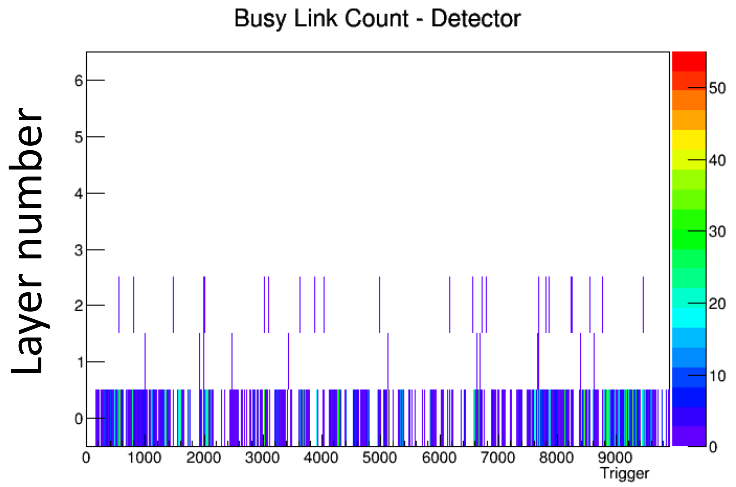

BUSY vs trigger, 10us strobe

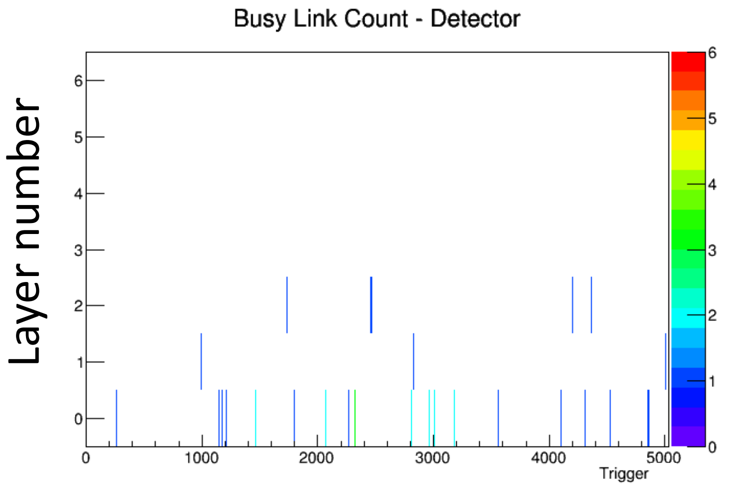

BUSY vs trigger, 20us strobe

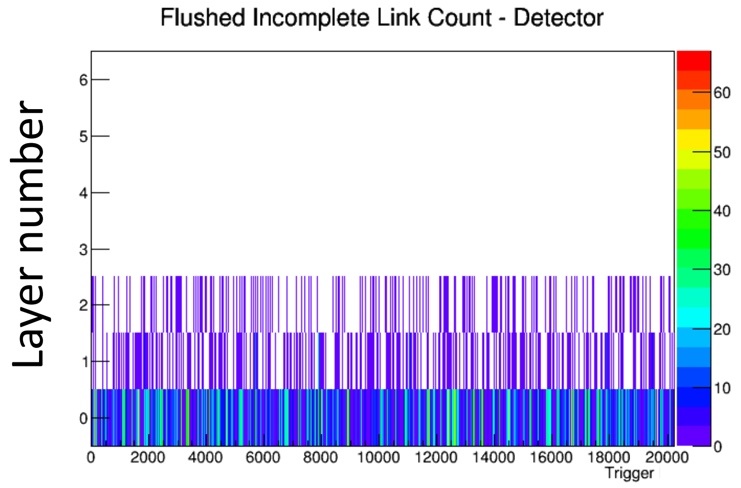

FLUSH INCOMPL. vs trigger, 5 us strobe

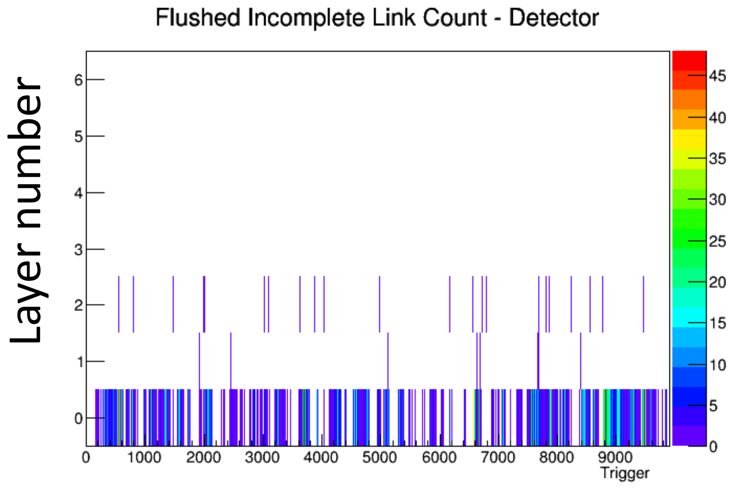

FLUSH INCOMPL. vs trigger, 10us strobe

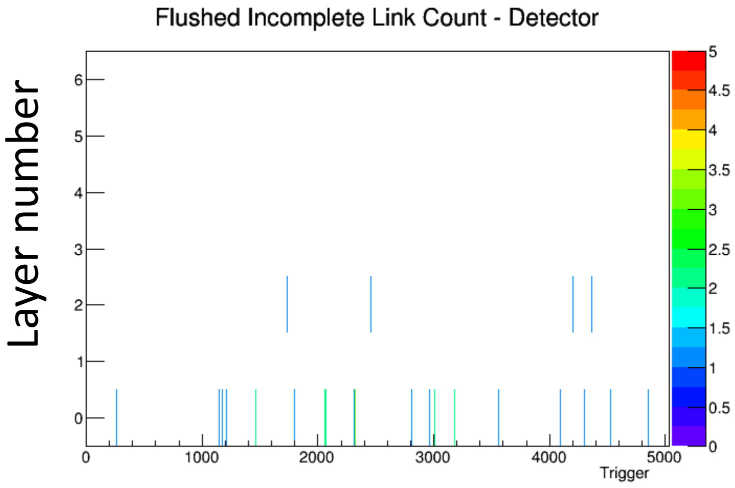

FLUSH INCOMPL. vs trigger, 20us strobe

PbPb simulations (continuous mode, 200kHz) – Results

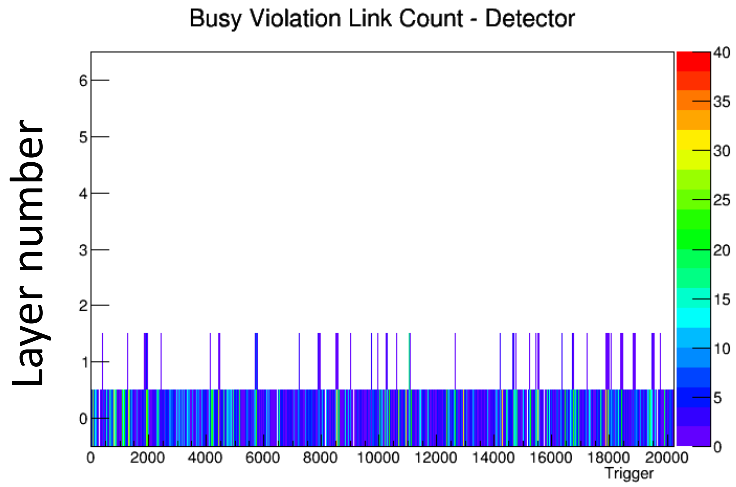

BUSYV vs trigger, 5us strobe

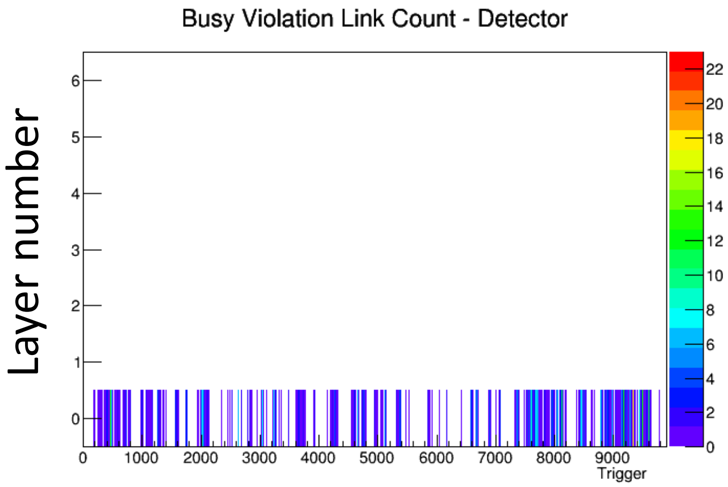

BUSYV vs trigger, 10us strobe

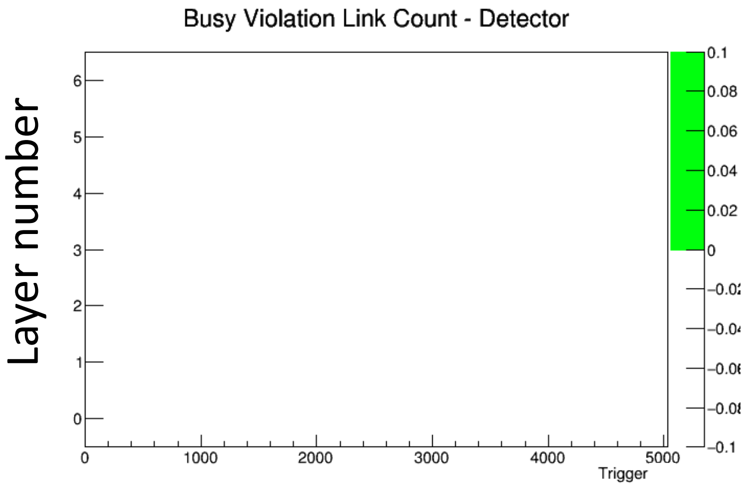

BUSYV vs trigger, 20us strobe

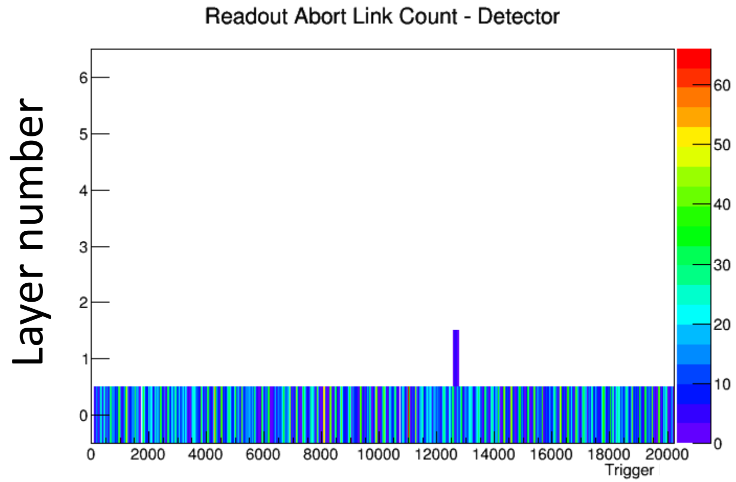

RO ABORT vs trigger, 5 us strobe

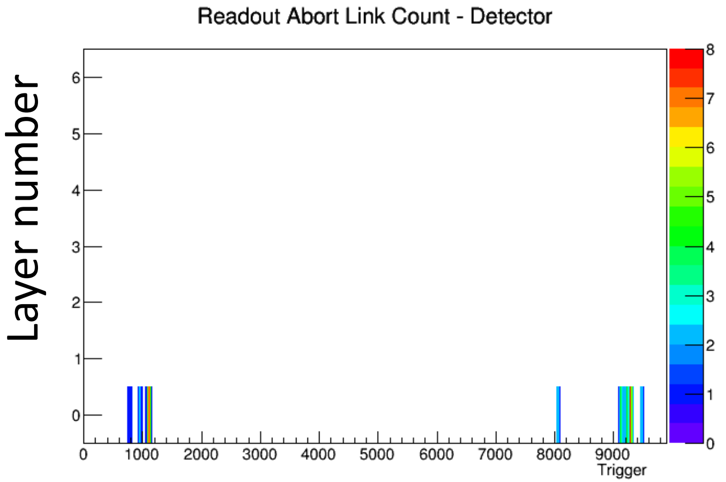

RO ABORT vs trigger, 10us strobe

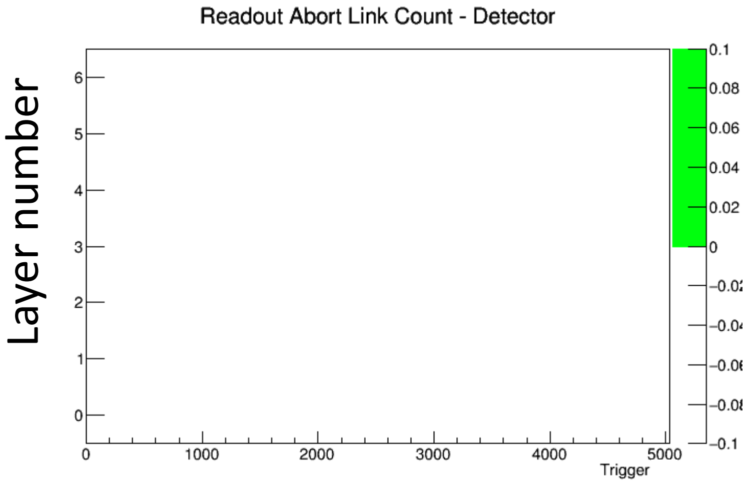

RO ABORT. vs trigger, 20us strobe

pp simulations – Event generation

A pool of discrete MC events for proton-proton were generated using the itsuTestBench in AliRoot:

- Using Pythia for random particle generation, and GEANT4 for tracking and detector response
  - Process: kPyMb
  - Tune: 14 (kPythia8Tune_Monash2013)
  - EnergyCMS: 5500
  - Momentum range: 0 to 999999
  - pT range: 0 to 1000
  - Theta range: 0 to 180 degrees
  - Y (rapidity) range: -2.5 to 2.5
- Random events were picked from the pool using a uniform distribution, and fed to the detector in SystemC
- Random time between events is picked from an exponential distribution, with lambda = 1 / avg_event_rate (time)

Rationale for high event-rate pp simulations

There is interest in probing production of rare heavy nuclei, such as helium/anti-helium, in proton-proton collisions

- To have a realistic chance of observing this, ALICE needs to run at higher interaction rates, maybe over 1MHz
- ALICE upgrade is designed for 200 kHz pp, ITS for 400 kHz.
- To investigate the viability of the upgraded ITS at interaction rates up to several MHz, simulations were run using the SystemC model

Readout Electronics Crates

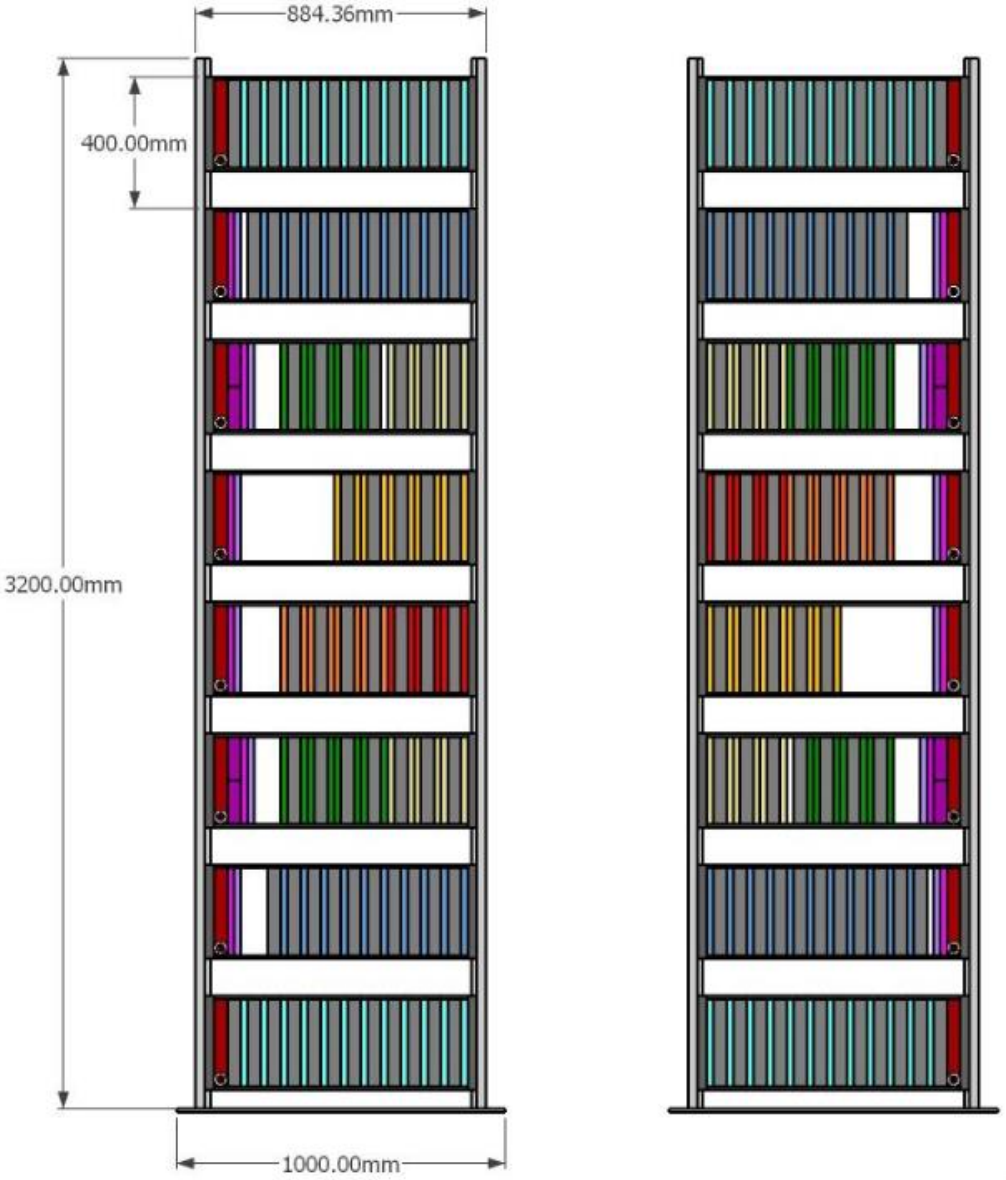

- Layer 0 RUs
- Layer 1 RUs
- Layer 2 RUs
- Layer 3 RUs
- Layer 4 RUs
- Layer 5 RUs
- Layer 6 RUs
- Power Boards (2PU each)
- Trigger splitters
- Busy Units
- Bias / CAN Bus breackout

Disclaimers

- OB not simulated
- Noise not included in simulations (very low, on the order of 1E-8)
- Number of simulated events is pretty low
  - Due to time constraints before PRR