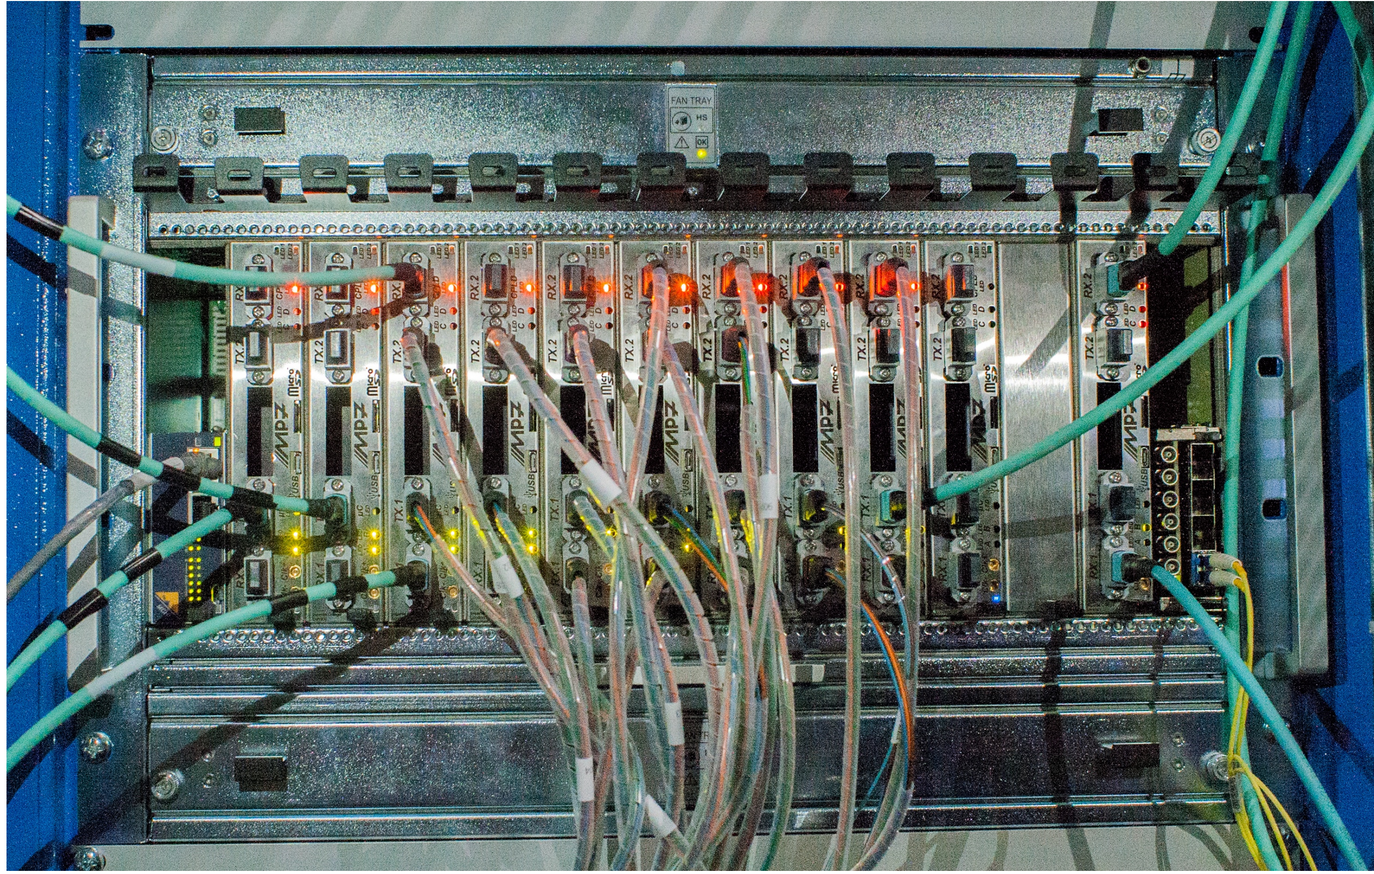

# FULL CHAIN DEMONSTRATOR

Tom James, Imperial College London

TMTT: L1 Tracking Review

08/Dec/2016

# SYSTEM ARCHITECTURE - THE TRACK FINDING PROCESSOR (TFP) - RECAP

- **Track Finding Processor (TFP)** boards receive links from adjacent detector (DTC) octants, corresponding to a single processing octant
- Each TFP processes 1/8 in φ & 1/tmp in time
- **No further duplication** or sharing between regions is required downstream
- System is compartmentalised
  - **One processing octant becomes the demonstrator slice unit**

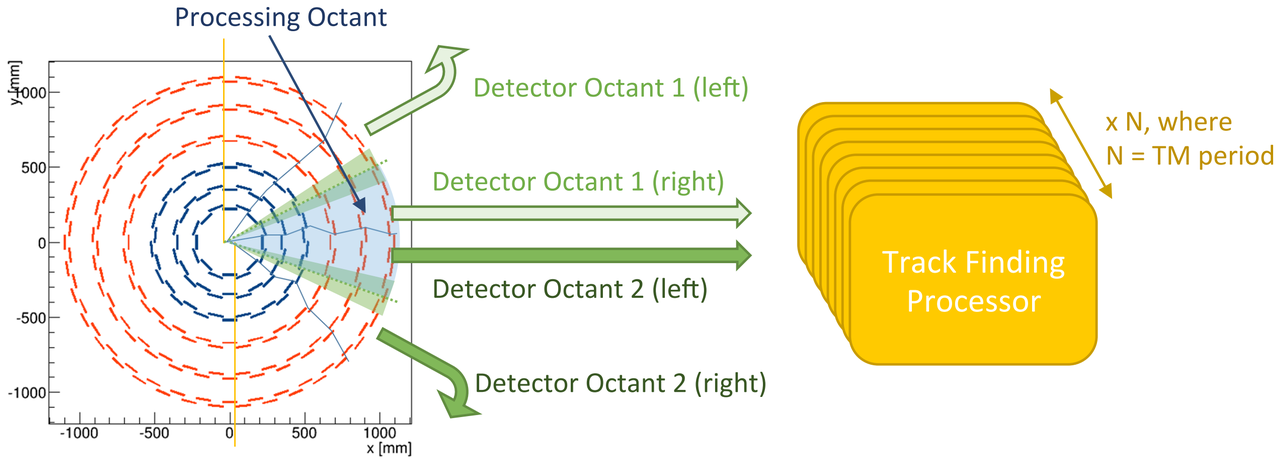

# SYSTEM ARCHITECTURE - MAPPING TO DEMONSTRATOR SLICE

- TFP is **internally divided into logical elements**, each on separate boards (MP7s)
- Simplifies division of labour and algorithm development/testing
- Presently available **FPGA resources** is **not necessarily a limit to the scale/performance** of algorithms we want to implement
- Can extrapolate FPGA resources to those of a future processing card, allowing demonstration of what could be final system performance with currently in hand technology

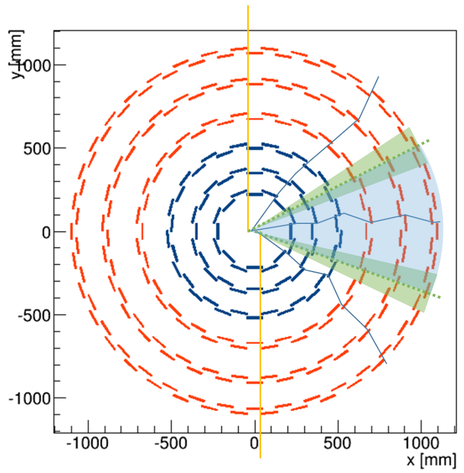


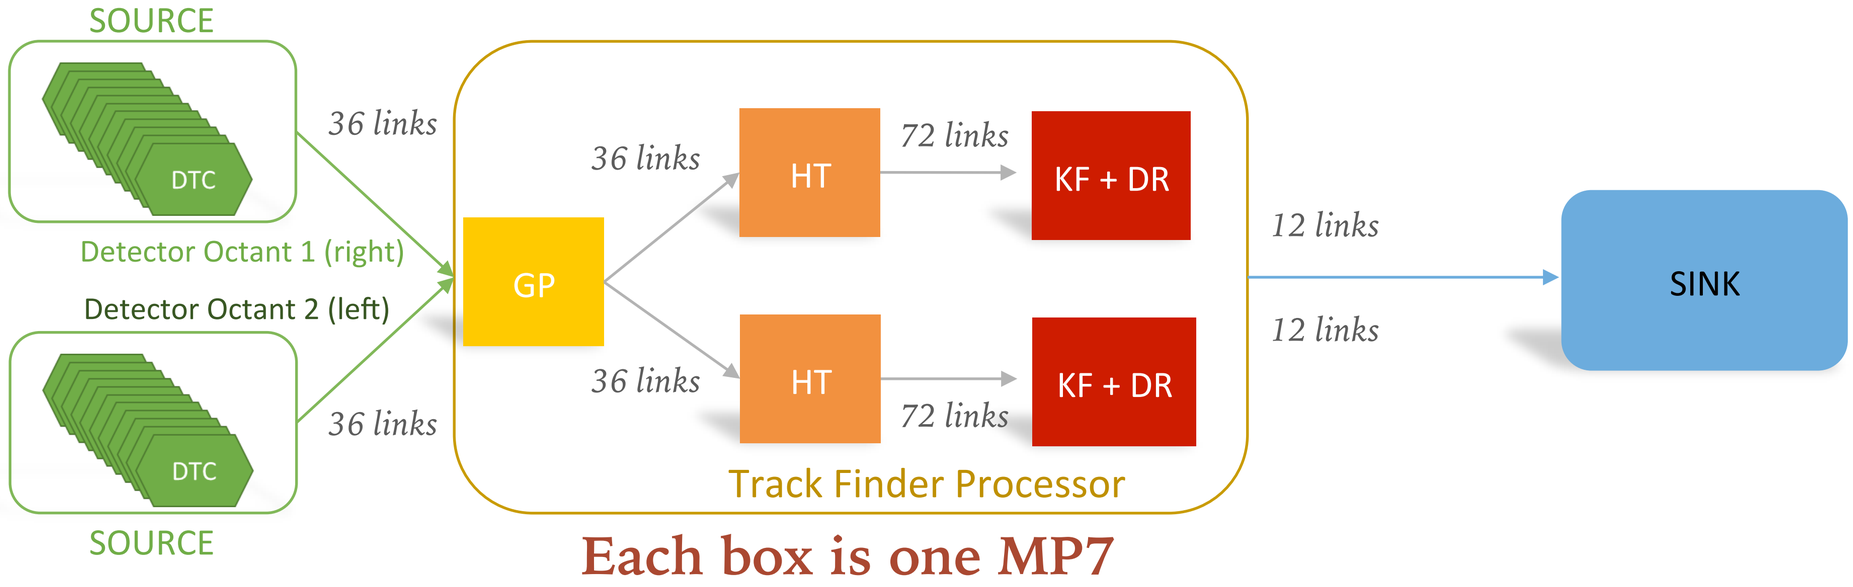


**Each box is one MP7**

# DEMONSTRATOR HARDWARE - THE MASTER PROCESSOR 7 (MP7)

- The MP7 is a generic high-performance data-stream processing double width AMC card, developed for use in the Phase I calorimeter trigger upgrade
- Equipped with a Xilinx Virtex-7 690-T FPGA
- 12 Avago Technologies MiniPOD optical transmitters/receivers
  - Each provide 12 optical links running at up to 10.3 Gbps
  - Total optical bandwidth 0.74 Tbps each way

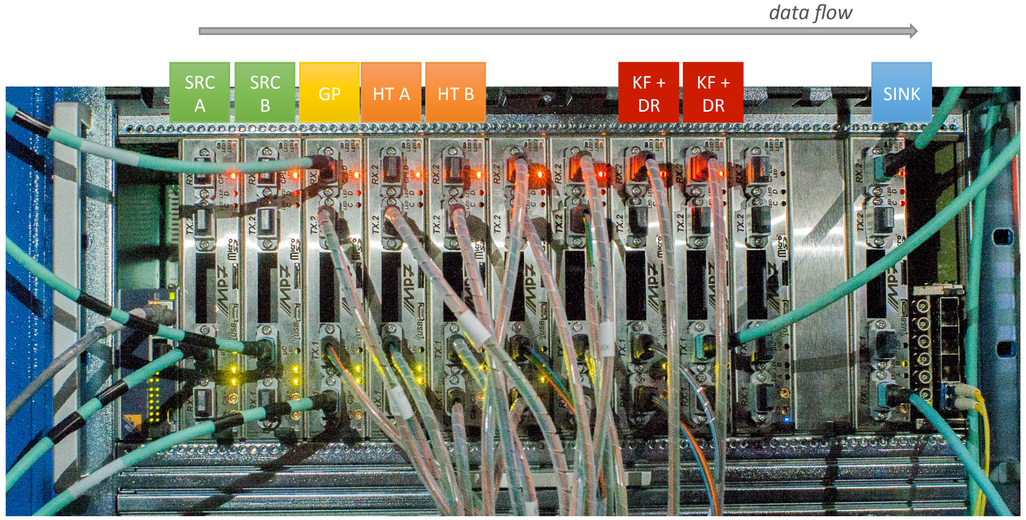


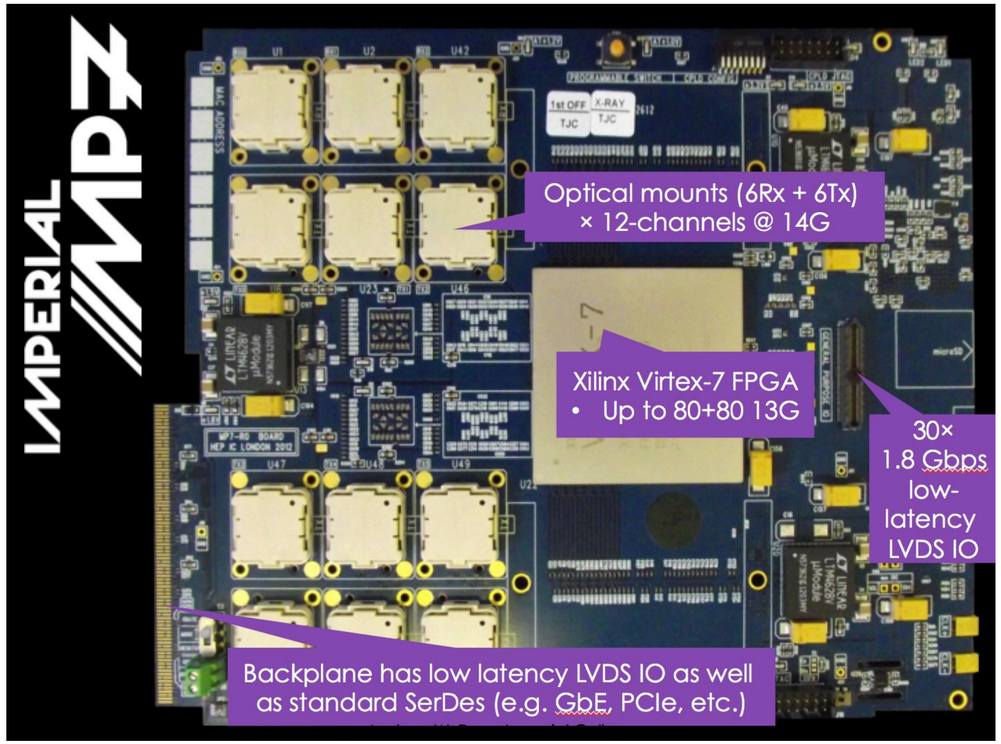

# DEMONSTRATOR HARDWARE - THE MASTER PROCESSOR 7 (MP7)

- Comes with **well supported firmware and software infrastructure**
- Firmware already provided for core tasks such as transceiver buffering, I/O formatting & external communication and configuration
  - Simplifies demonstration process greatly -> **can focus on the algorithms**
- **Using well understood hardware means that details of timing and synchronisation were not an issue - we benefit from a great deal of work previously done**

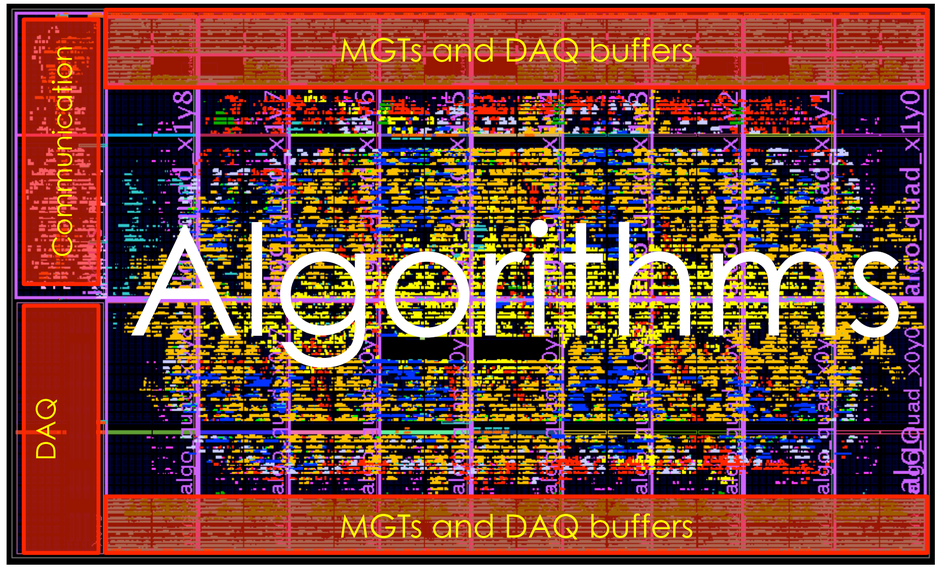


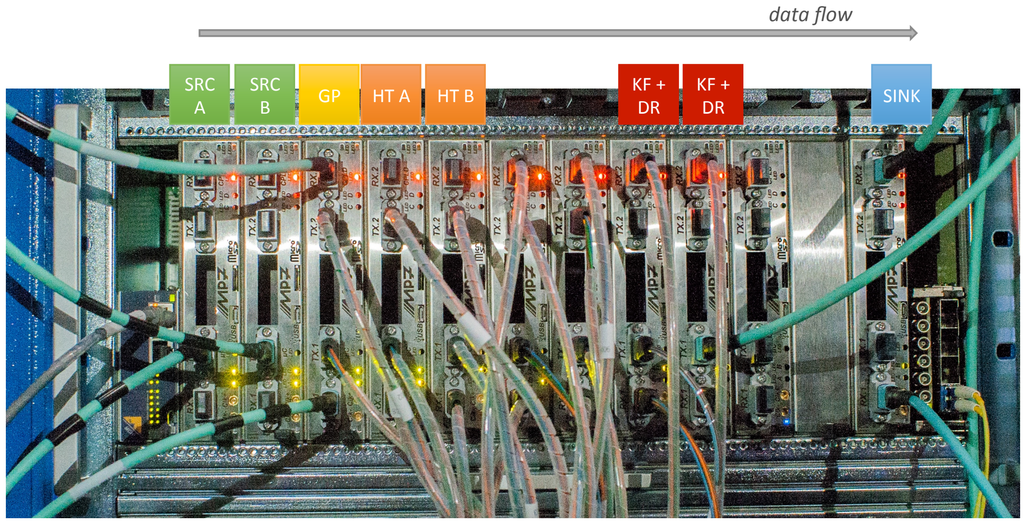

# DEMONSTRATOR HARDWARE

- Location: **Tracker Integration Facility (TIF) B186, CERN**
- **CERN Blue rack -** Turbine, 3-phase power, air deflector, **water cooling/heat exchangers**
- **Schroff** MicroTCA crate powered by external PowerOne 48V PSU & Vadatech power modules
- Equipped with NAT-**MCH (Gigabit Ethernet communication via backplane) & AMC13 (synchronisation, timing & control)**
- PowerEdge R620 CMS rack PC
- 11 **MP7-XE**'s installed in Schroff crate
  - 5 for algorithm demonstration (one TFP)
  - 3 as large buffers for source & sink
  - 3 spares/backup
- Boards daisy-chained with optical fibres to meet required demonstrator configuration

# DEMONSTRATOR OVERVIEW - DATA TAKING

- 8 *daisy-chained* MP7 boards
- Five boards emulate one Track Finder Processor
- Processes Monte Carlo stubs for **any one octant in φ, all of η at once**
- We take take data for all eight octants to **generate hardware results for entire tracker**

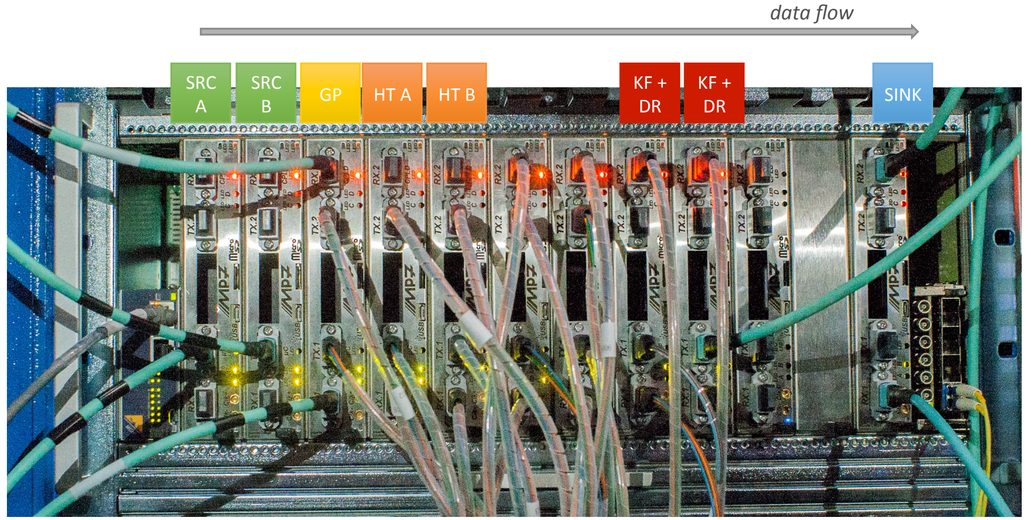


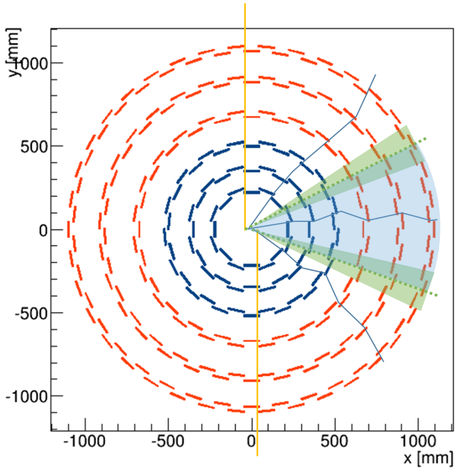


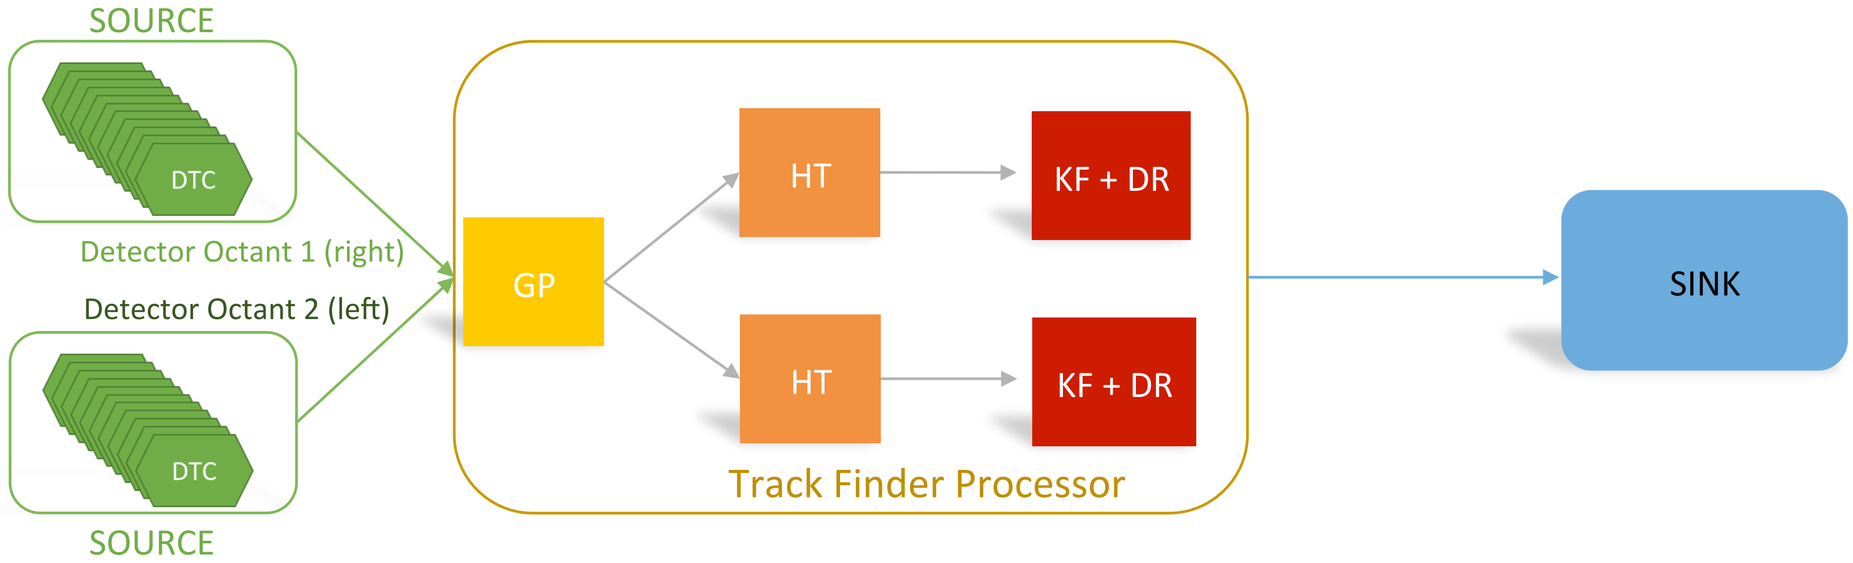

# DEMONSTRATOR OVERVIEW - SOURCE & SINK

- The **source** represents up to 36 virtual DTCs, covering a φ-octant in both z+ and z-
- Where a virtual DTC is a **time multiplexed stream of data as if it were coming from a real DTC**
- **Source** stores **~30 events** for playback
- The **sink** stores output of **~30 events**

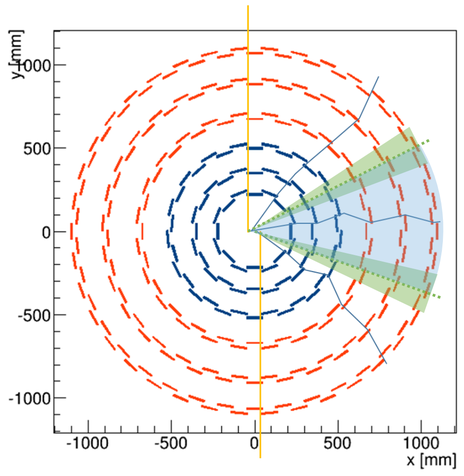


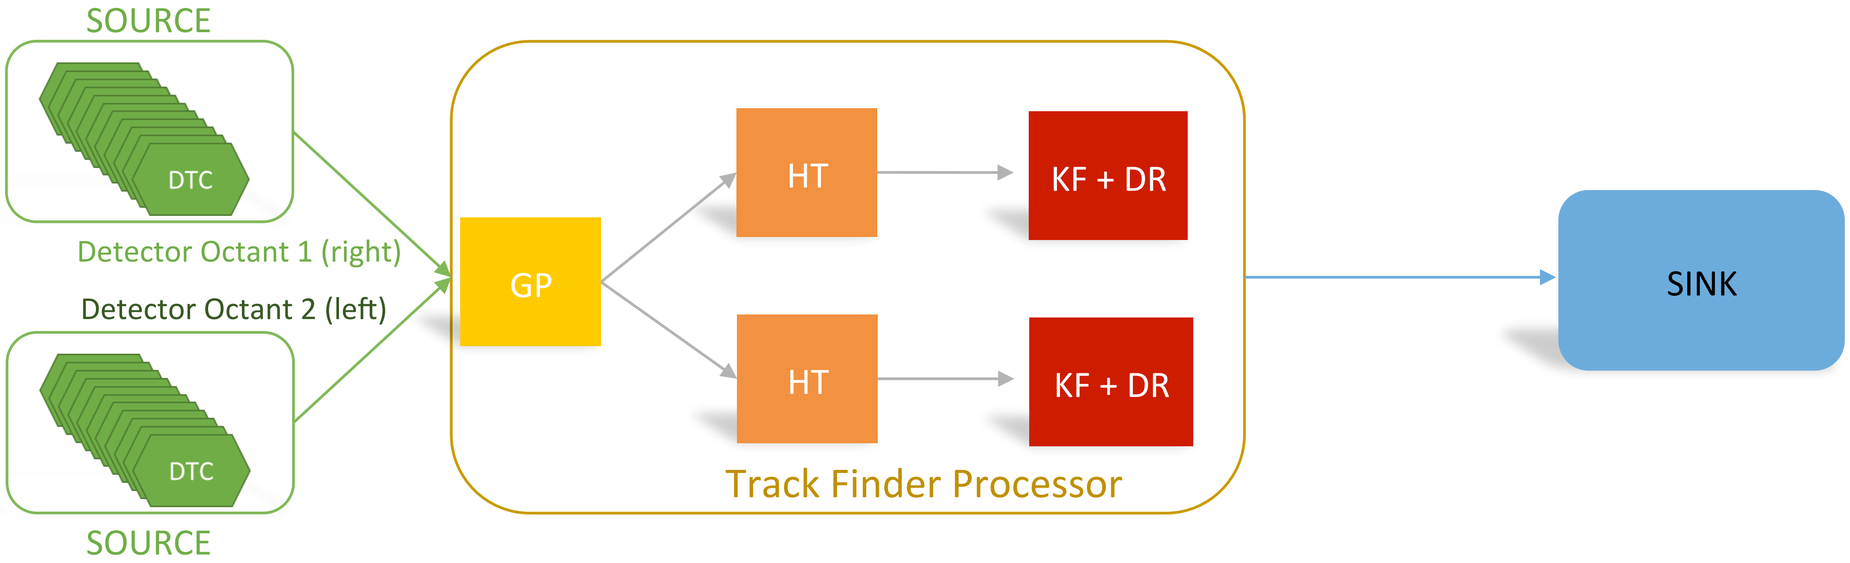

# DEMONSTRATOR OVERVIEW - SOURCE & SINK

- **Same fw for both source & sink**
- **72 Big Buffers** 16k deep - Dual port **RAMS**
- Acts like a **FIFO**
- 16k deep rams -> 8k 32 bit half-stubs, as 2 bits must be used for data valid & strobe
- Read/Write via **IPBus**

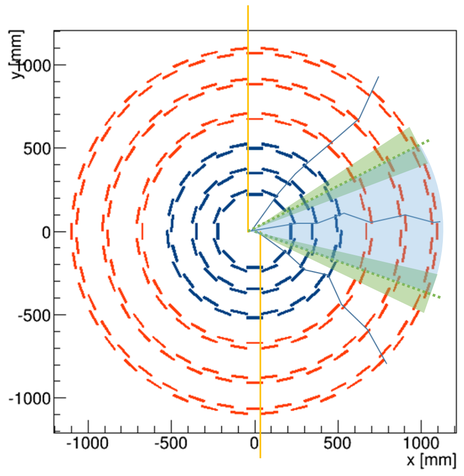


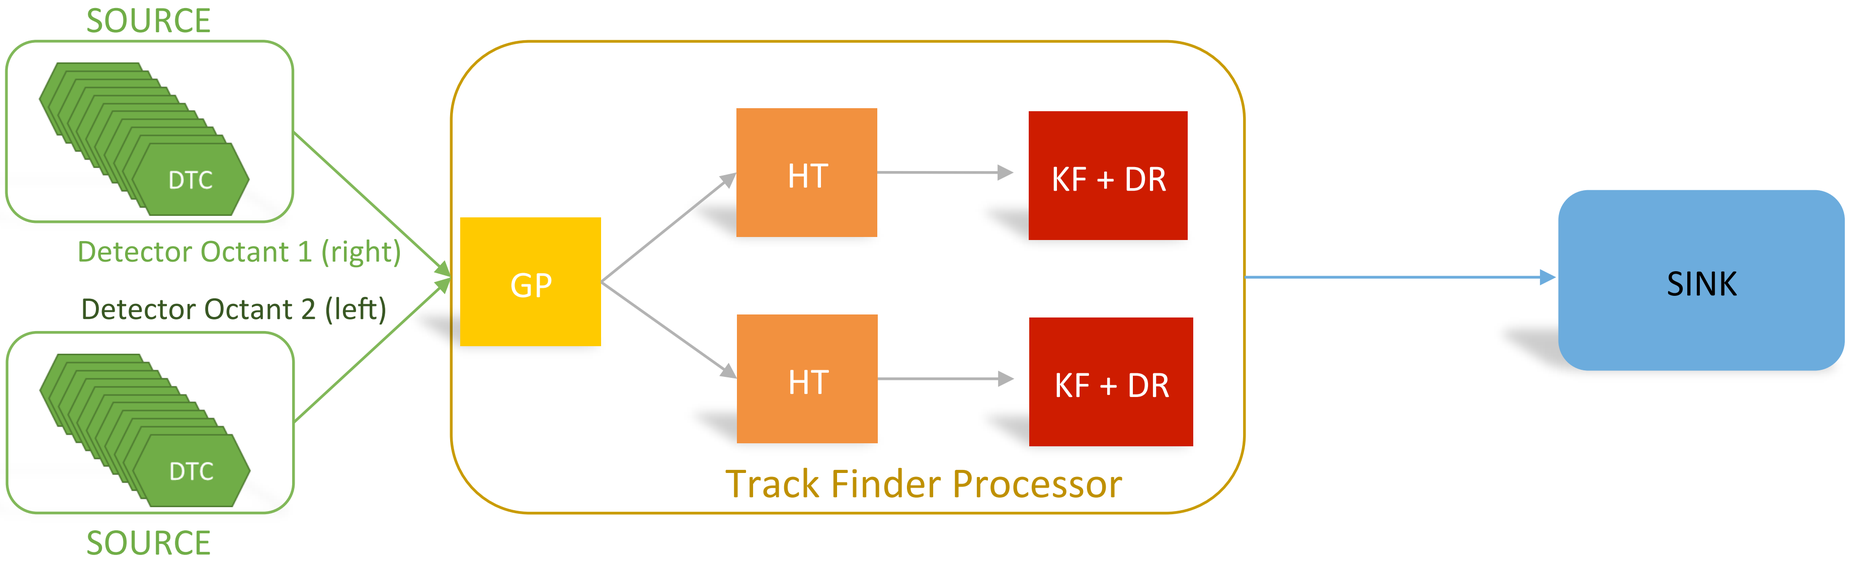

# DEMONSTRATOR OVERVIEW

We compare hardware output **directly** with cmssw simulation software -> **can measure performance directly with hardware**

- **Objective - To run standard physics samples through a hardware demonstrator to ensure that expected performance, as seen in simulation results, is realistic**
- Full MC events passed through hardware, and **tracks found are compared with those found by our CMSSW simulation/emulation software**

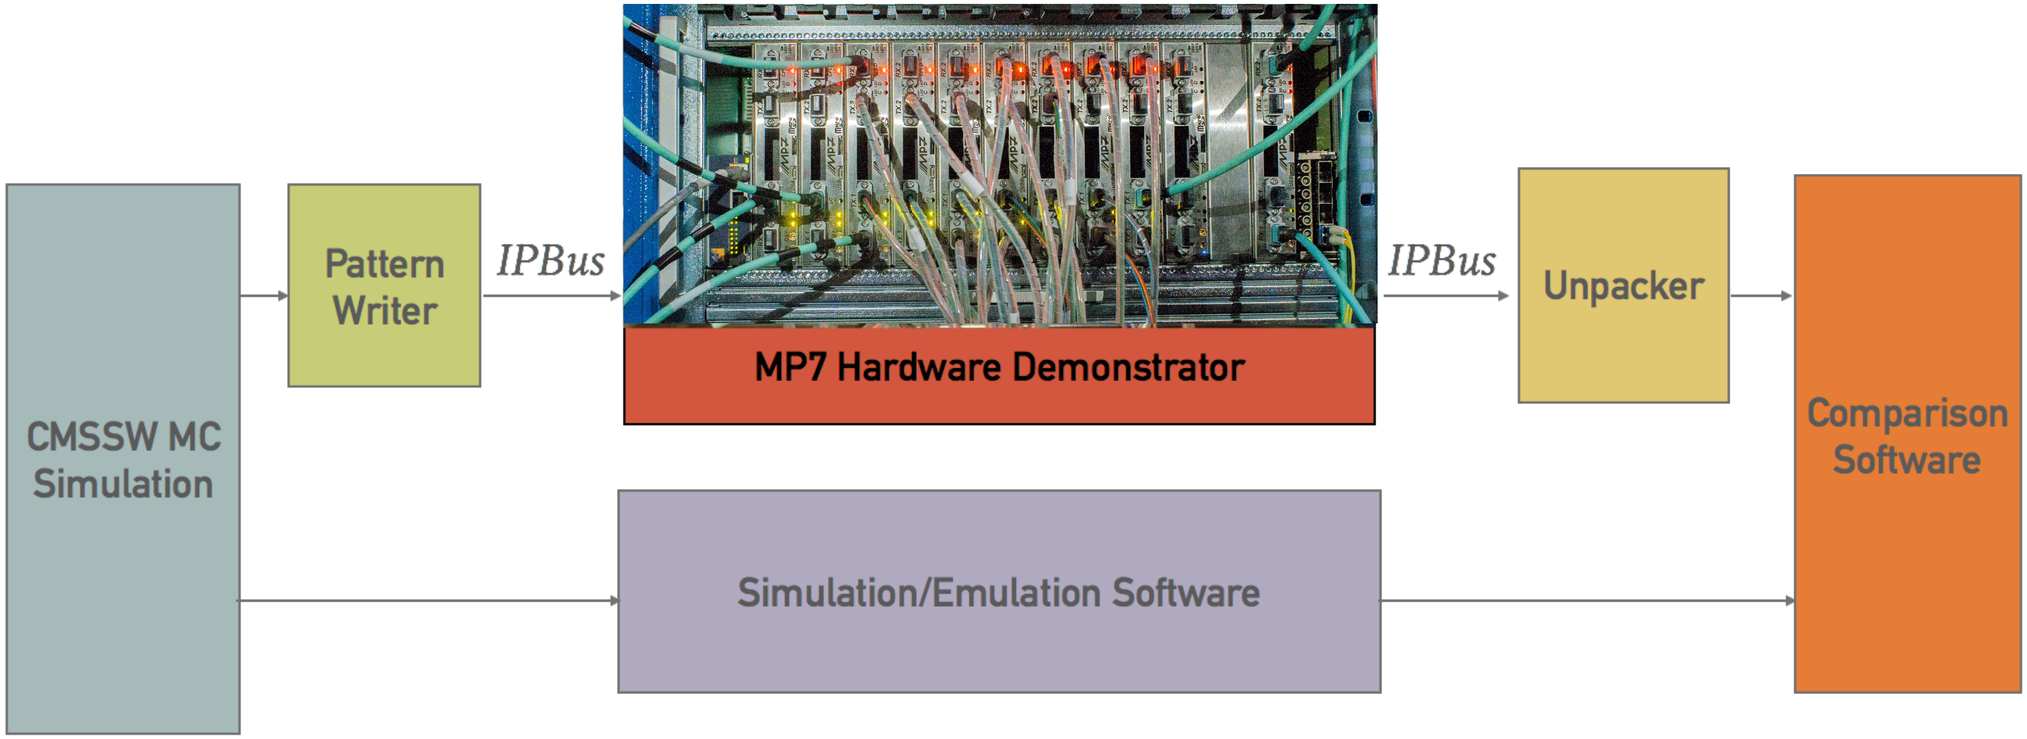

# DEMONSTRATOR RESULTS - TTBAR + 200 PU (1800 EVENTS)

- Highly efficient at finding TTBar inclusive tracks
- Demonstrator matching CMSSW simulation well

| | Efficiency (%) | Av Rate | Matched tracks (%) |
|---|---|---|---|
| hw | 94.5 | 76.5 | 98.7 |
| sw | 94.8 | 79.4 | |

All results use official matching criteria

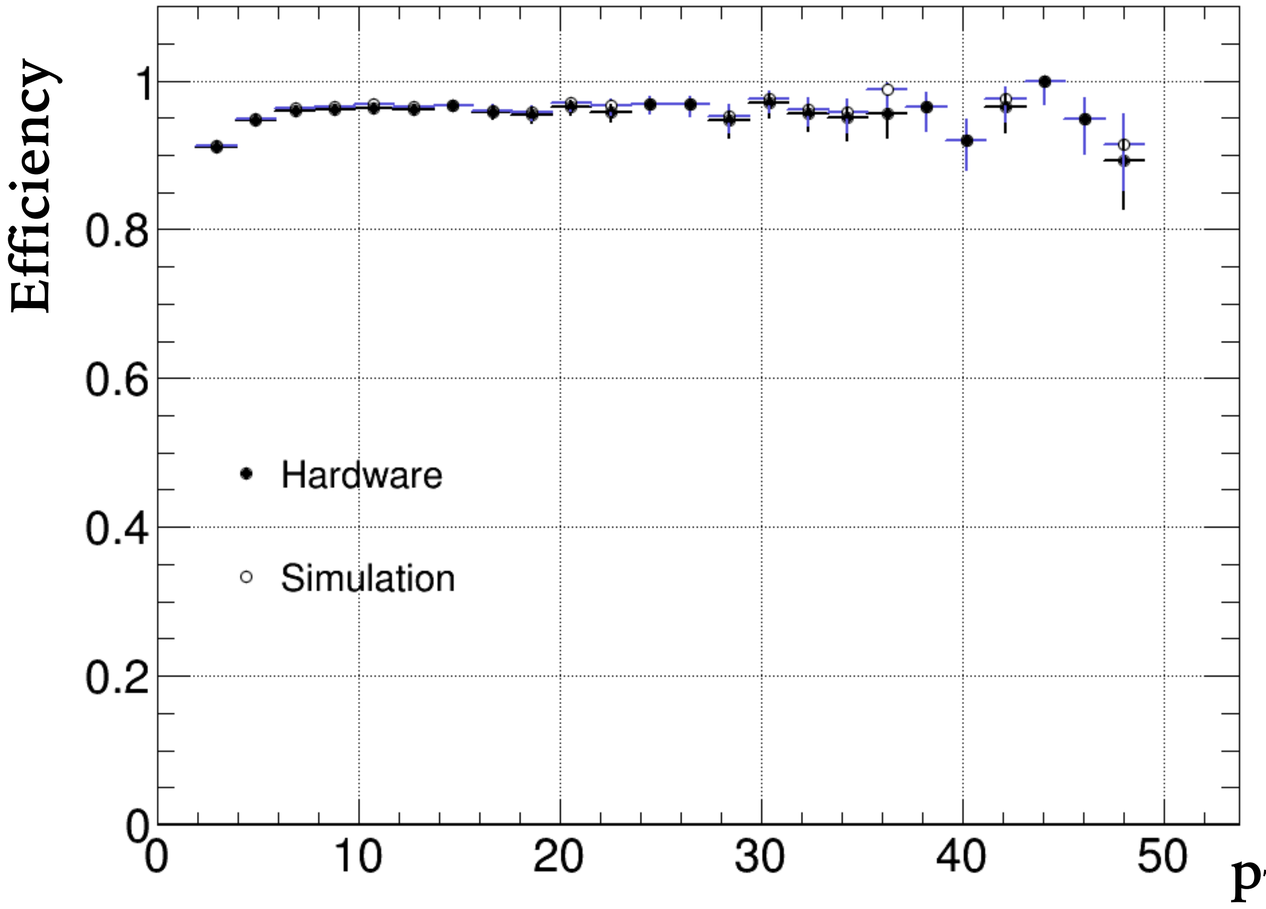


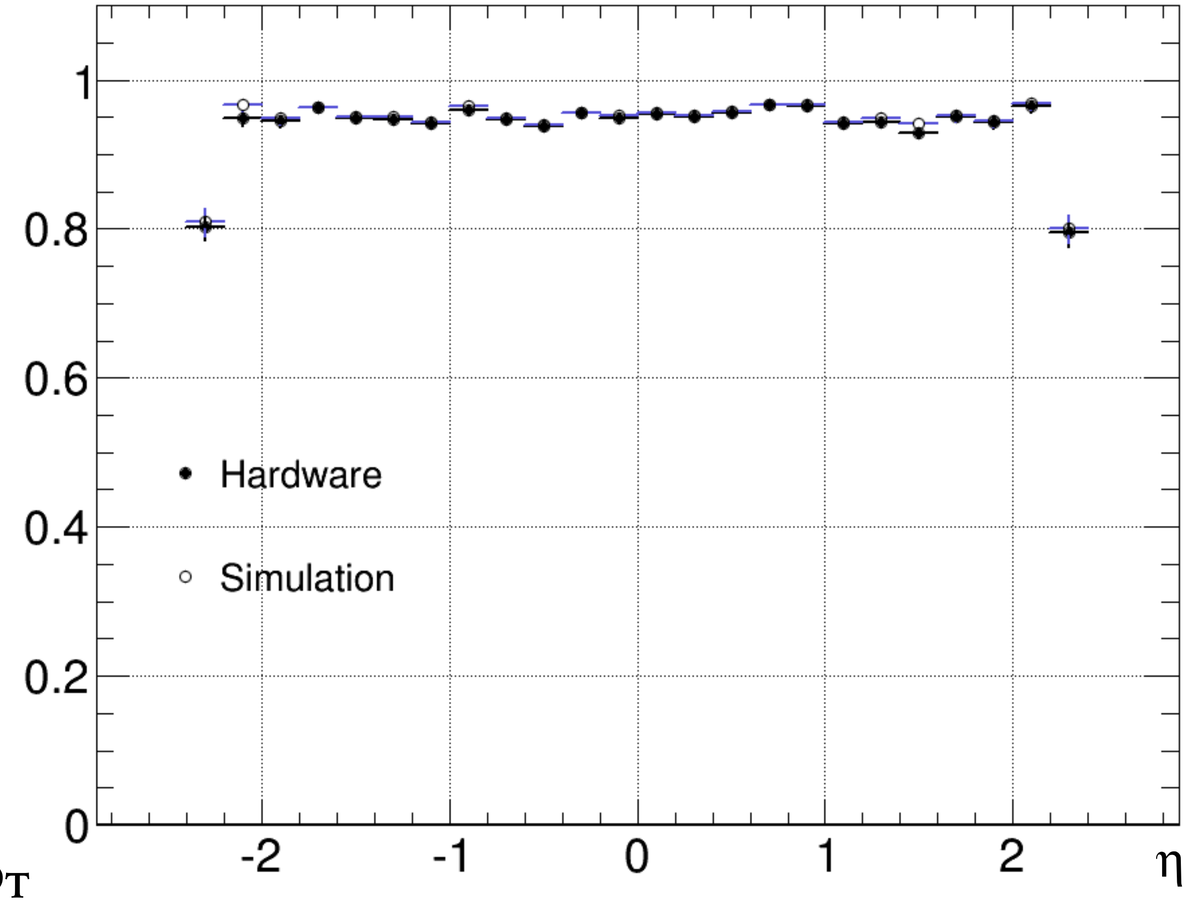

# DEMONSTRATOR RESULTS - MUONS 8-100 GEV + 200 PU (1100 EVENTS)

| | Efficiency (%) | Matched tracks (%) |
|---|---|---|
| hw | 97.1 | 99.2 |
| sw | 97.1 | |

- Highly efficient at finding single muons in 200 PU
- Efficiency matches CMSSW simulation exactly

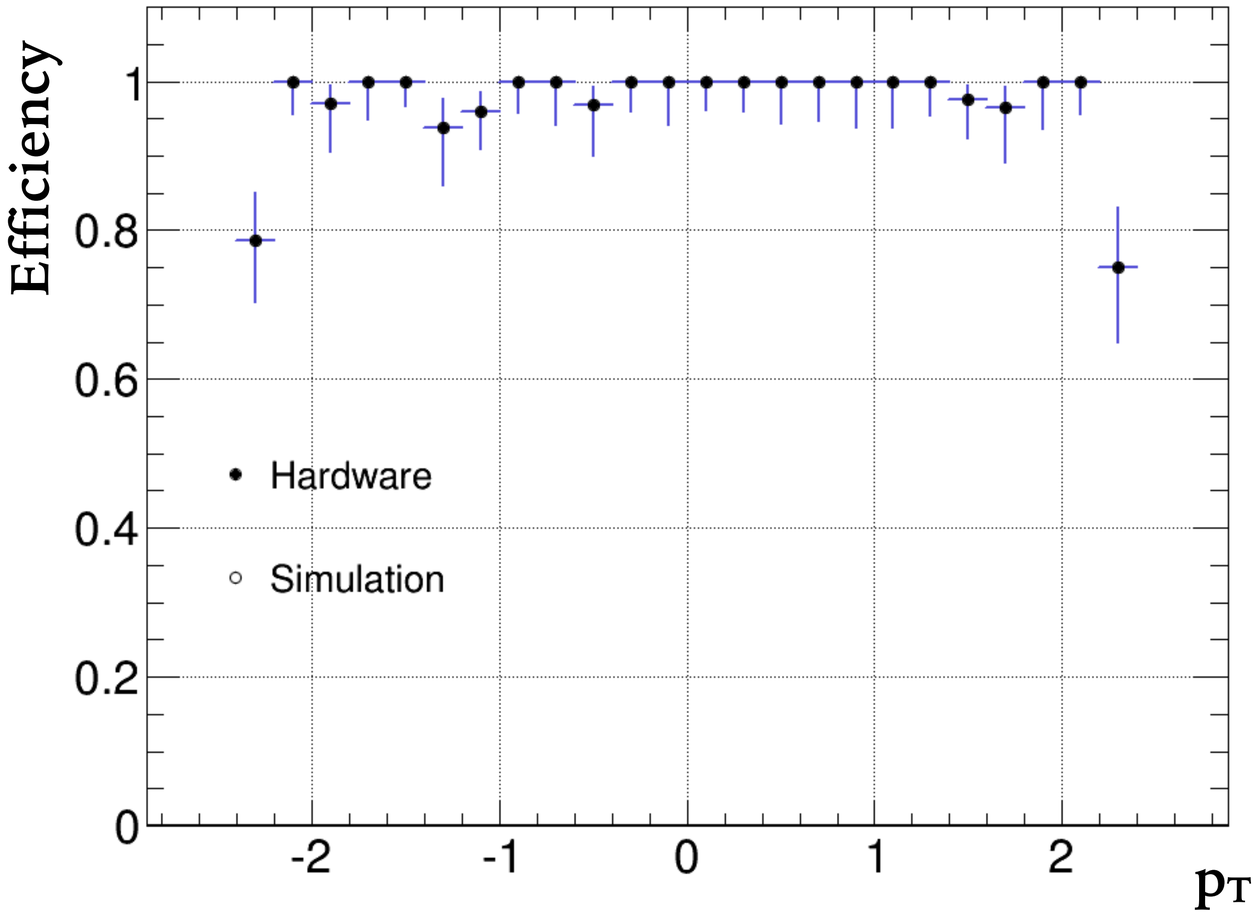


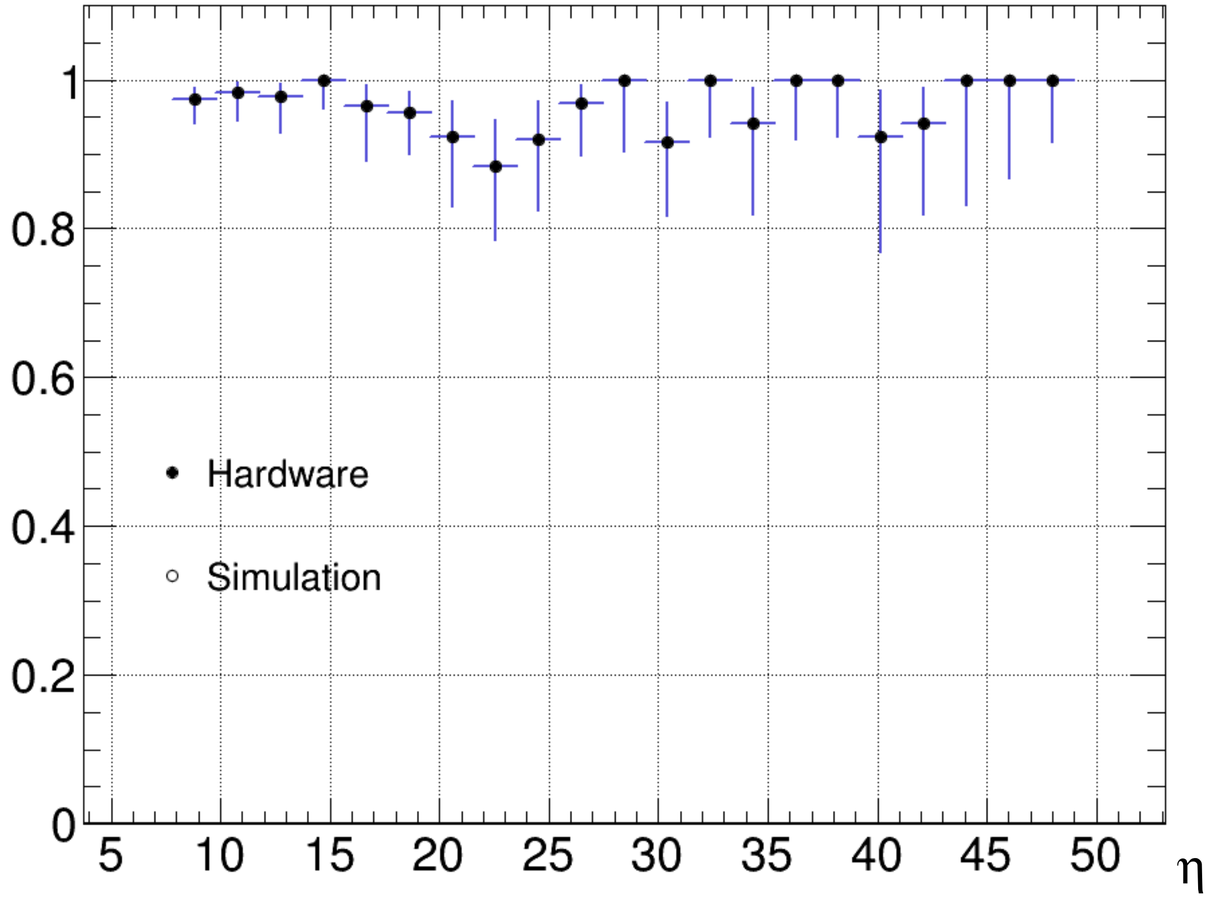

# DEMONSTRATOR RESULTS - MUONS 8-100 GEV + 200 PU (1100 EVENTS)

- Well matched efficiency in CMSSW and demonstrator allow for extrapolation of results to the higher statistics currently available in software

| | Efficiency (%) | Matched tracks (%) |
|---|---|---|
| hw | 97.1 | 99.2 |
| sw | 97.1 | |

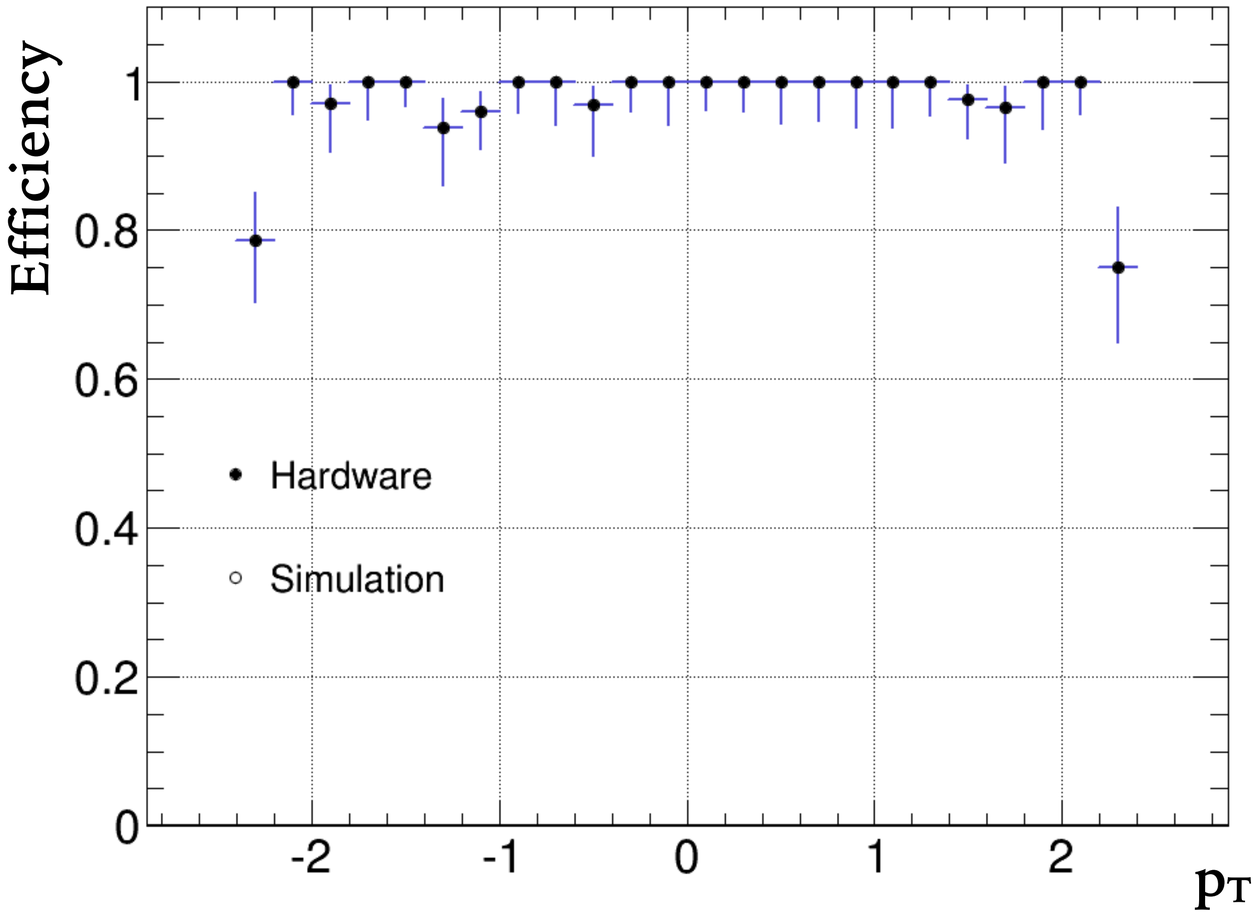


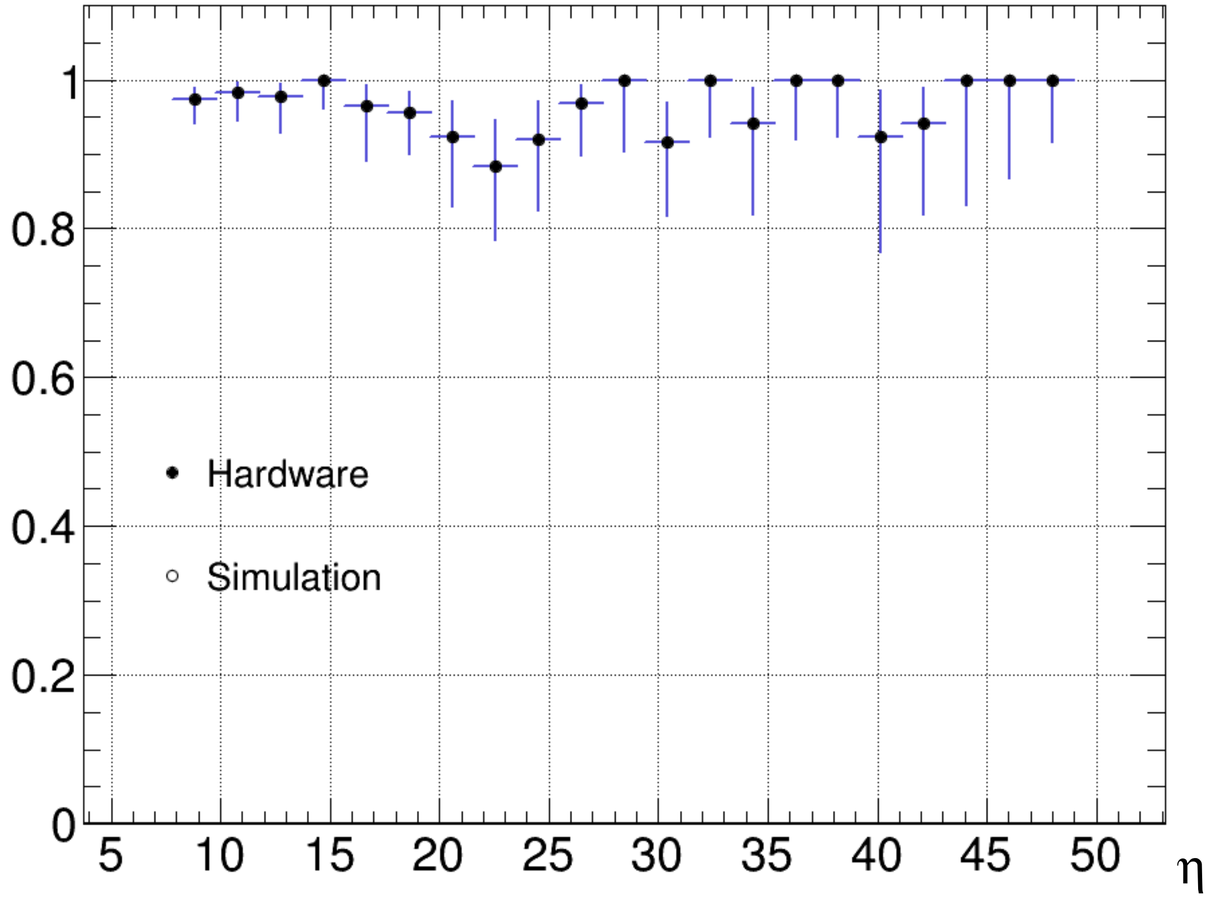

# DEMONSTRATOR RESULTS - ELECTRONS IN TTBAR + 200 PU (1800 EVENTS)

| | Efficiency (%) | Matched tracks (%) |
|---|---|---|
| hw | 81.4 | 98.7 |
| sw | 81.8 | |

- Particle-gun electron samples not available until yesterday
  - Using electrons in ttbar + 200 PU samples instead
- Performance matches CMSSW simulation

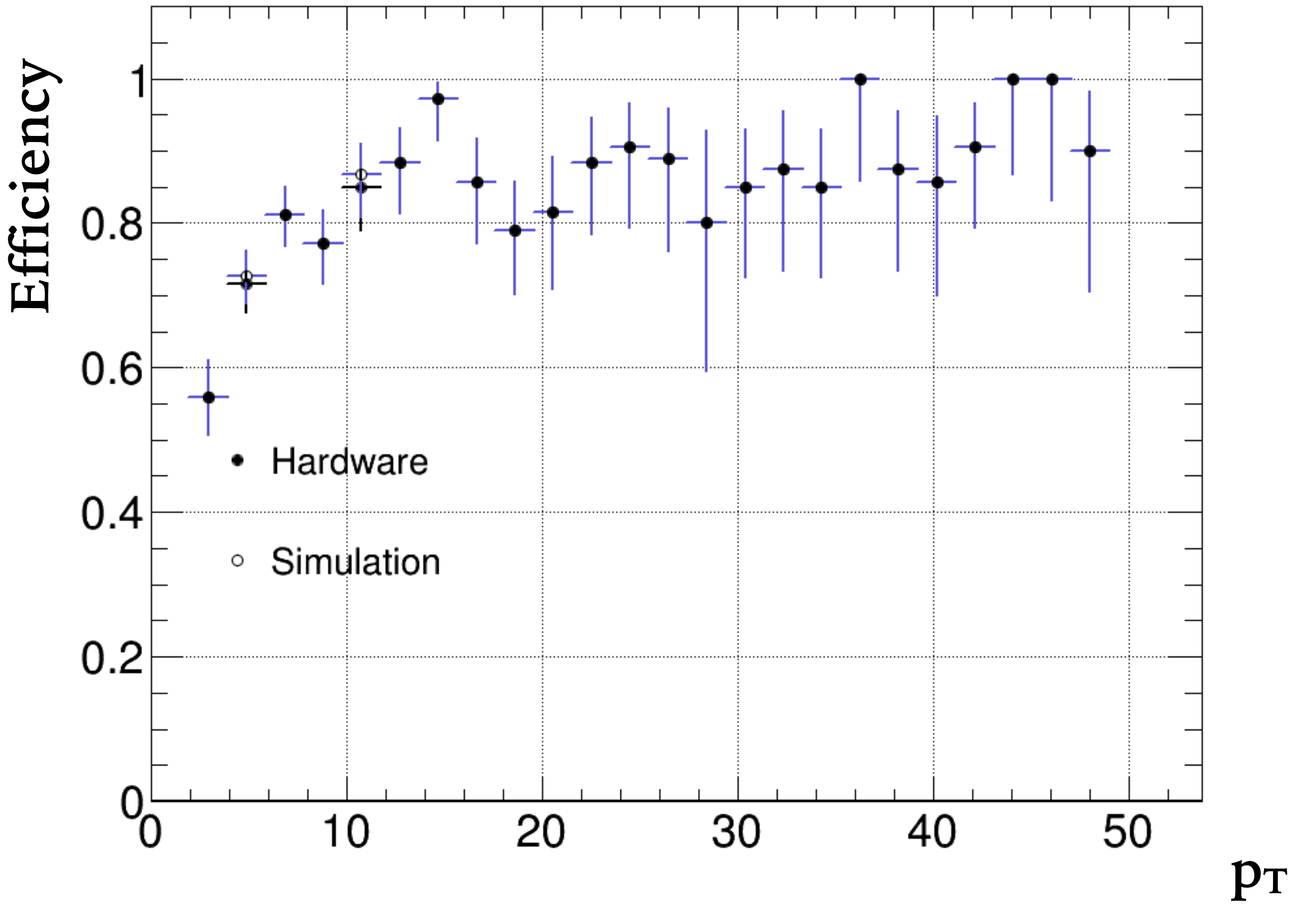

# DEMONSTRATOR RESULTS - DIGITISED RESOLUTIONS (TTBAR + 200PU)

- Resolution of track helix parameters good in simulation and hardware
- Realised very recently that demonstrator $z_0$ resolution was degraded by our choice of 12 and 10 bit encoding of r and z stub coordinates
- Simulation shows we can recover optimal resolution in demonstrator by using 2 more bits to encode:
  - r of stubs in the barrel
  - z in endcaps
  - Can trivially accommodate this change in demonstrator without degrading resolution of other helix parameters, but did not have time before review
  - Software configuration for demonstrator comparison plots use the smaller number of bits for better matching

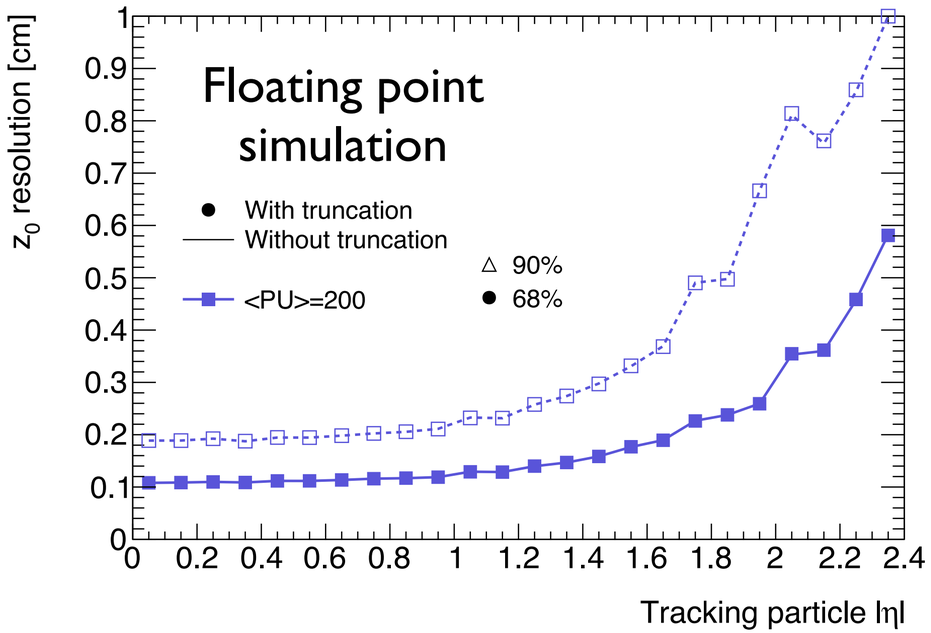


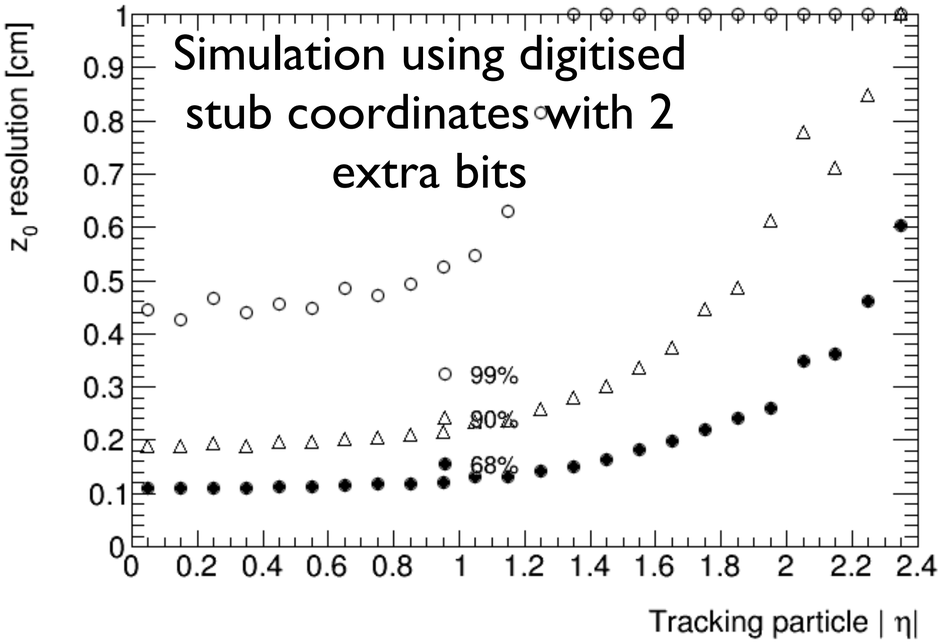

# DEMONSTRATOR RESULTS - TTBAR + 200PU RESOLUTIONS (1800 EVENTS)

- Excellent helix parameter resolutions measured in demonstrator

## $q/p_T$ resolution [1/GeV]

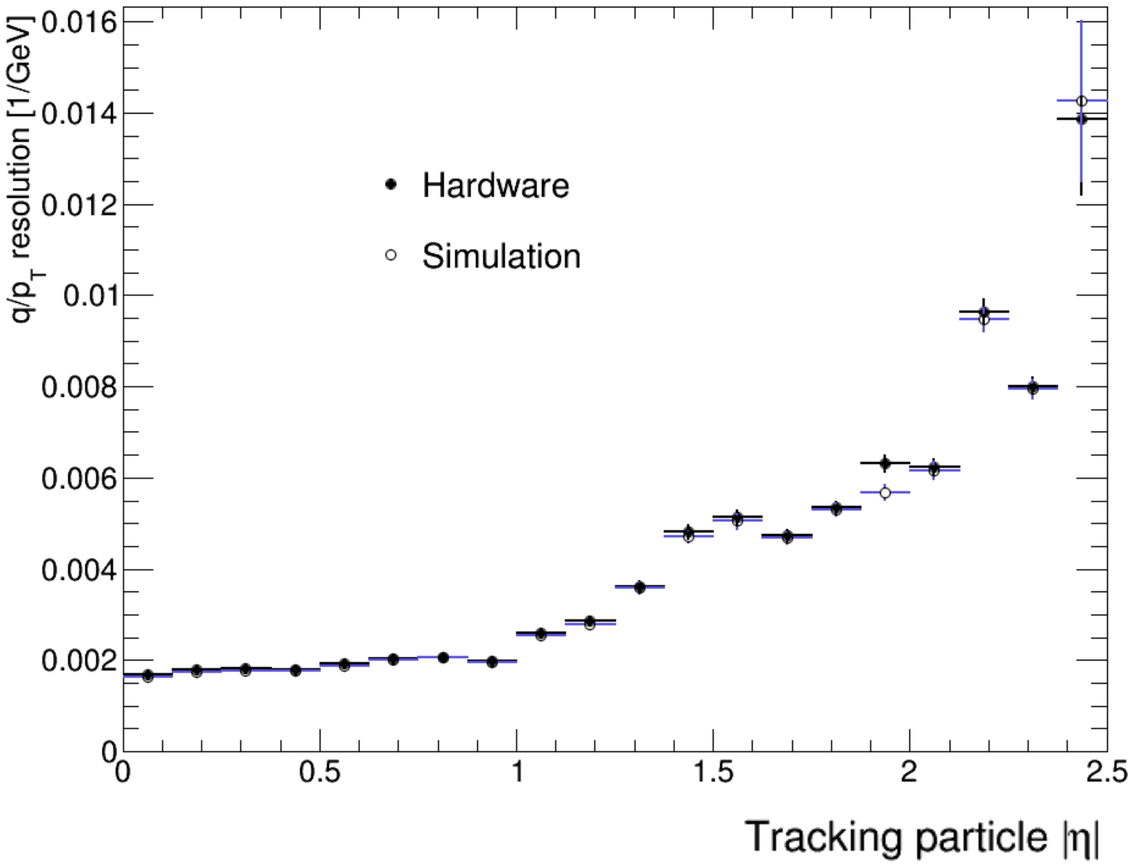


## $\varphi_0$ resolution [rad]

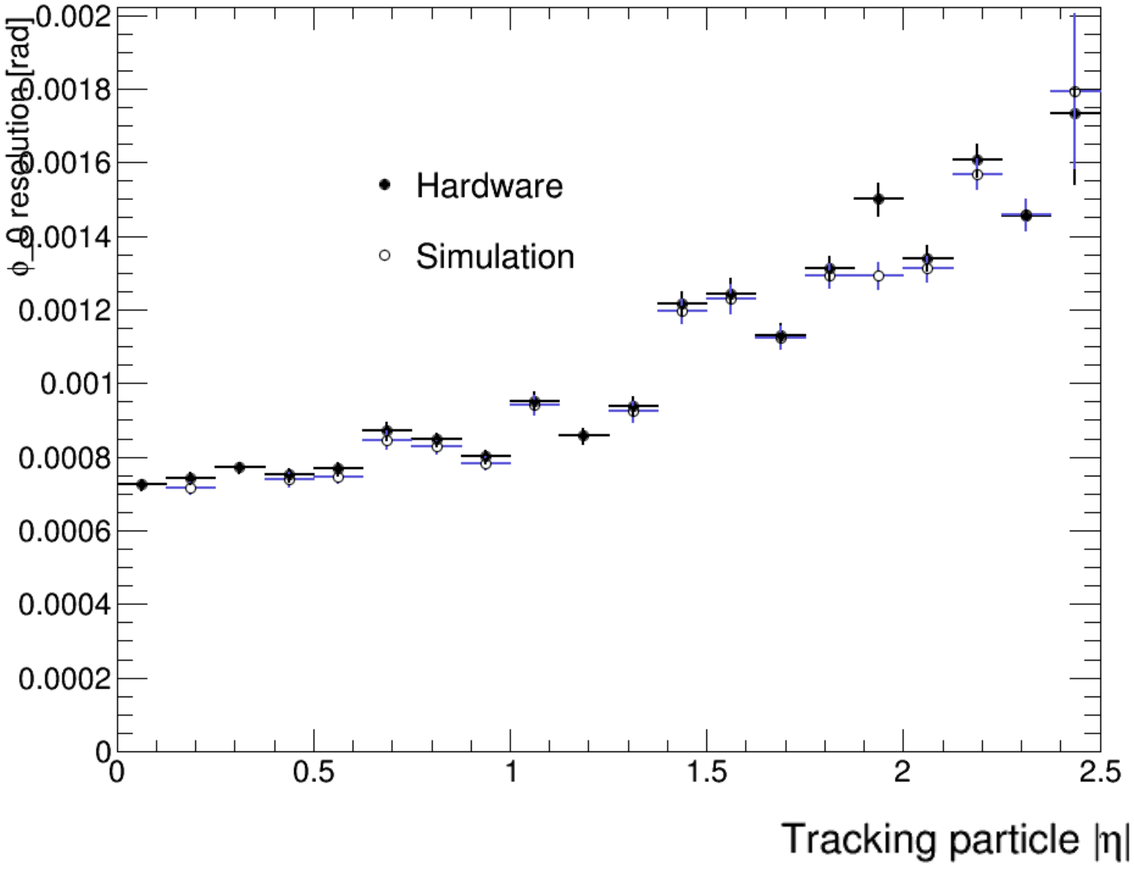

# DEMONSTRATOR RESULTS - TTBAR + 200PU RESOLUTIONS (1800 EVENTS)

➤ Excellent helix parameter resolutions measured in demonstrator

## $z_0$ resolution [cm]

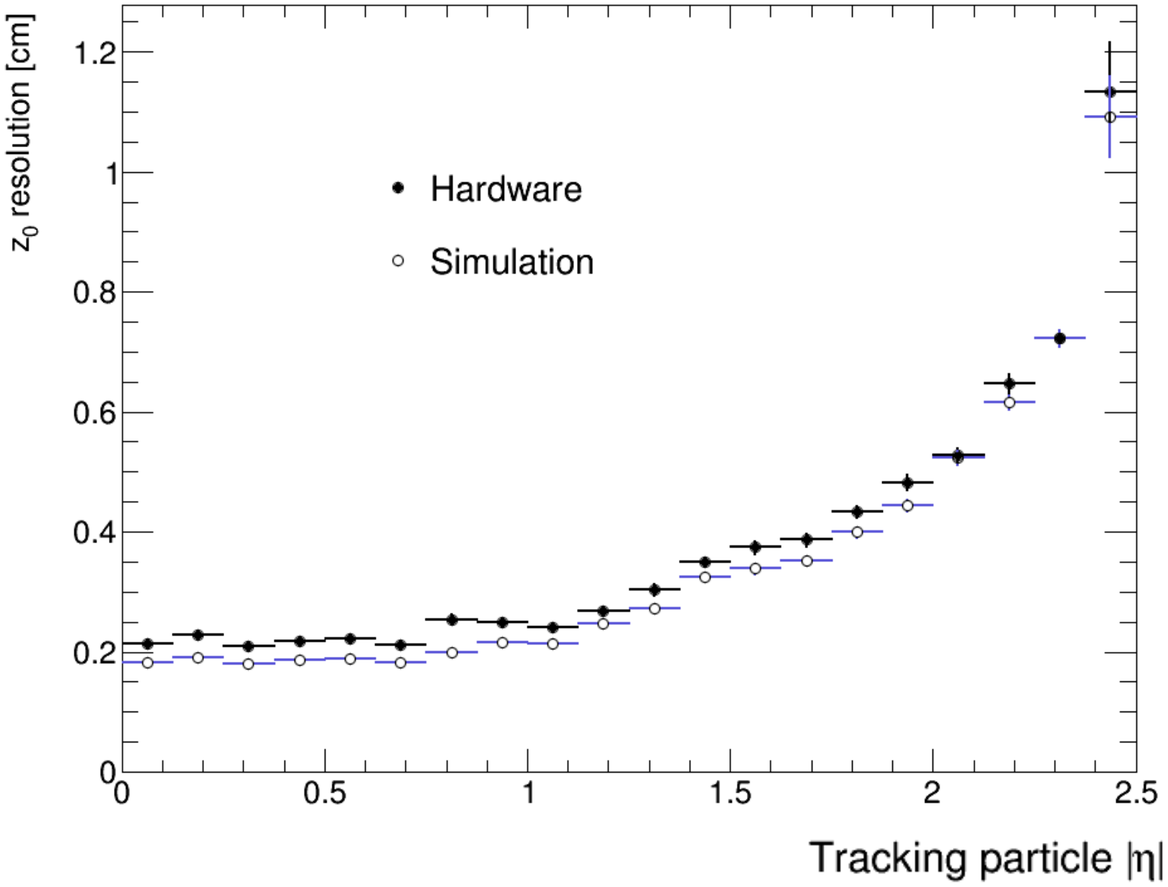


## η resolution

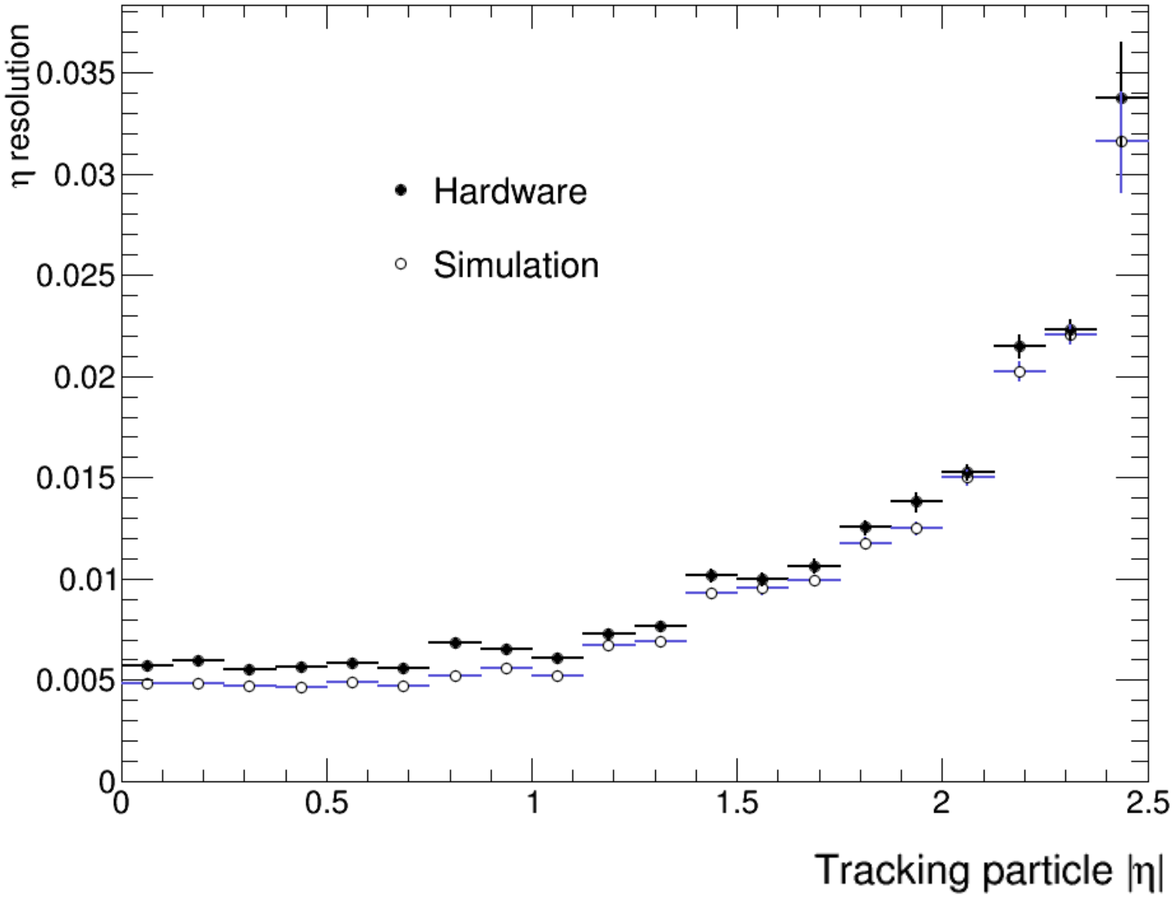

# DEMONSTRATOR RESULTS - RATE FOR TTBAR + 200PU

- Av. rate out tracks of the duplicate removal stage is measured in hw as ~76 tracks per event
- Duplicate removal recently integrated into processor chain and shows good results
- Small discrepancies in duplicate rate will be debugged over the next couple of weeks

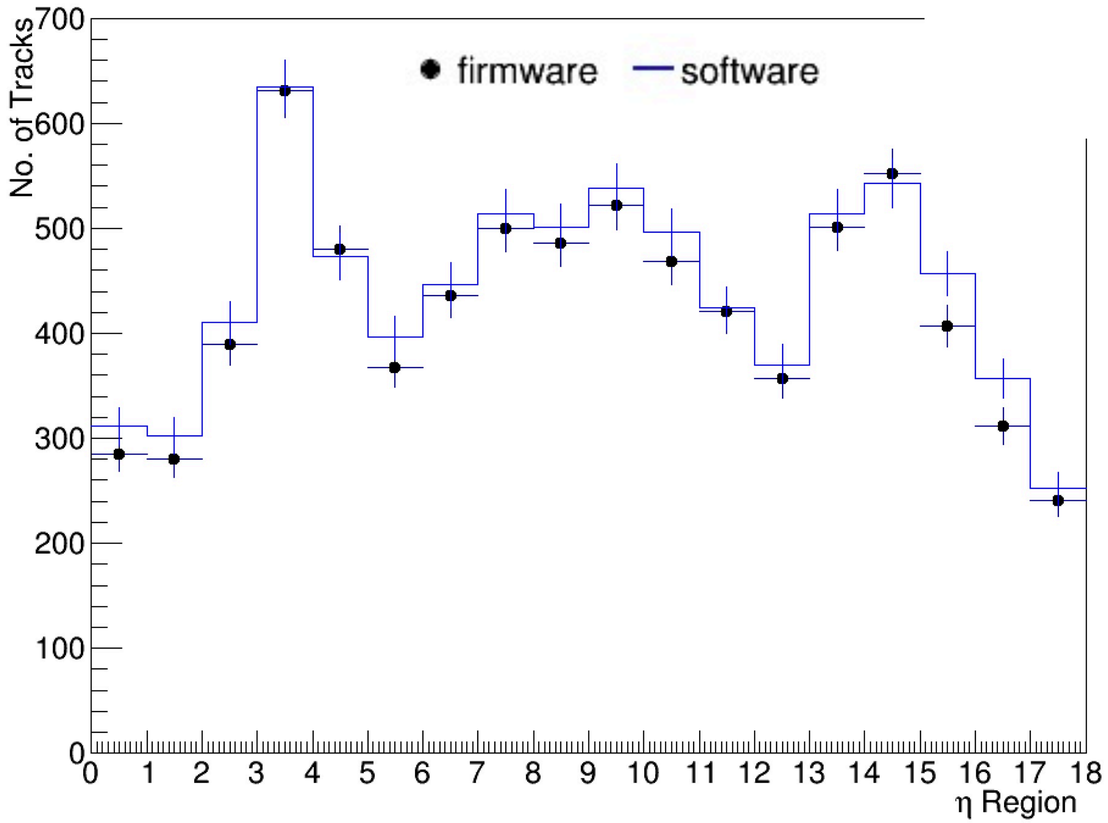


| TTbar + 200 PU | Av Rate |
|---|---|
| hw | 76.5 |
| sw | 79.4 |

# DEMONSTRATOR RESULTS - DEAD COOLING LOOP SCENARIO - TTBAR + 200 PU (900 EVENTS)

- As seen in CMSSW simulation, the demonstrator can also be configured to recover performance in a dead cooling loop scenario
- Online configuration of Hough Transform over ipbus all that is required

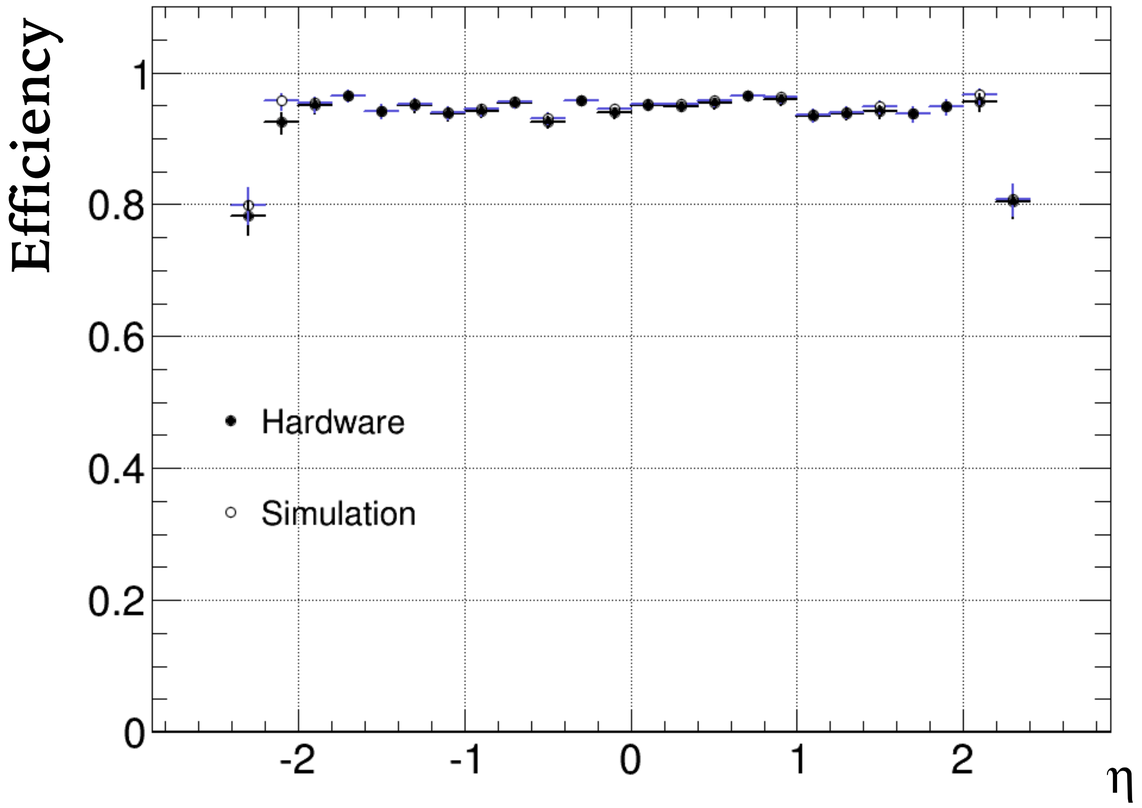


- In this dead cooling loop example, eta regions 5-8 have been configured to accept HT candidates with only 4 stubs
- Average efficiency is preserved at 94.6% in hardware

| | Efficiency (%) | Matched tracks (%) |
|---|---|---|
| hw | 94.2 | 98.4 |
| sw | 94.5 | |

# DEMONSTRATOR RESULTS - FPGA RESOURCES

- Although we are using 5 MP7's, to demonstrate one *Track Finding Processor,* the actual resource usage of the system is much smaller than we have available
- One can see that the GP+HT, and the KF+DR could each fit inside Ultrascale or Ultrascale+ generation chips

| per Tracker octant one TFP | LUTS | BRAM (36Kb) | DSPs |
|---|---|---|---|
| GP + HT | 412k | 1566 | 1560 |
| KF + DR | 382k | 1750 | 5040 |
| Virtex 7 690 | 433k | 1470 | 3600 |
| Kintex Ultrascale 115 | 663k | 2160 | 5520 |
| Virtex Ultrascale+ 11P | 1296k | 1970 | 9216 |

# DTC REQUIREMENTS

- TMTT requirements of the DTC FPGA
  - Conversion to global coordinates (48 bits)
  - Sorting data by event into N time multiplexed streams (demonstrator N = 36)
  - Duplicating data at our processing node boundaries, so no cross-node data flow downstream required

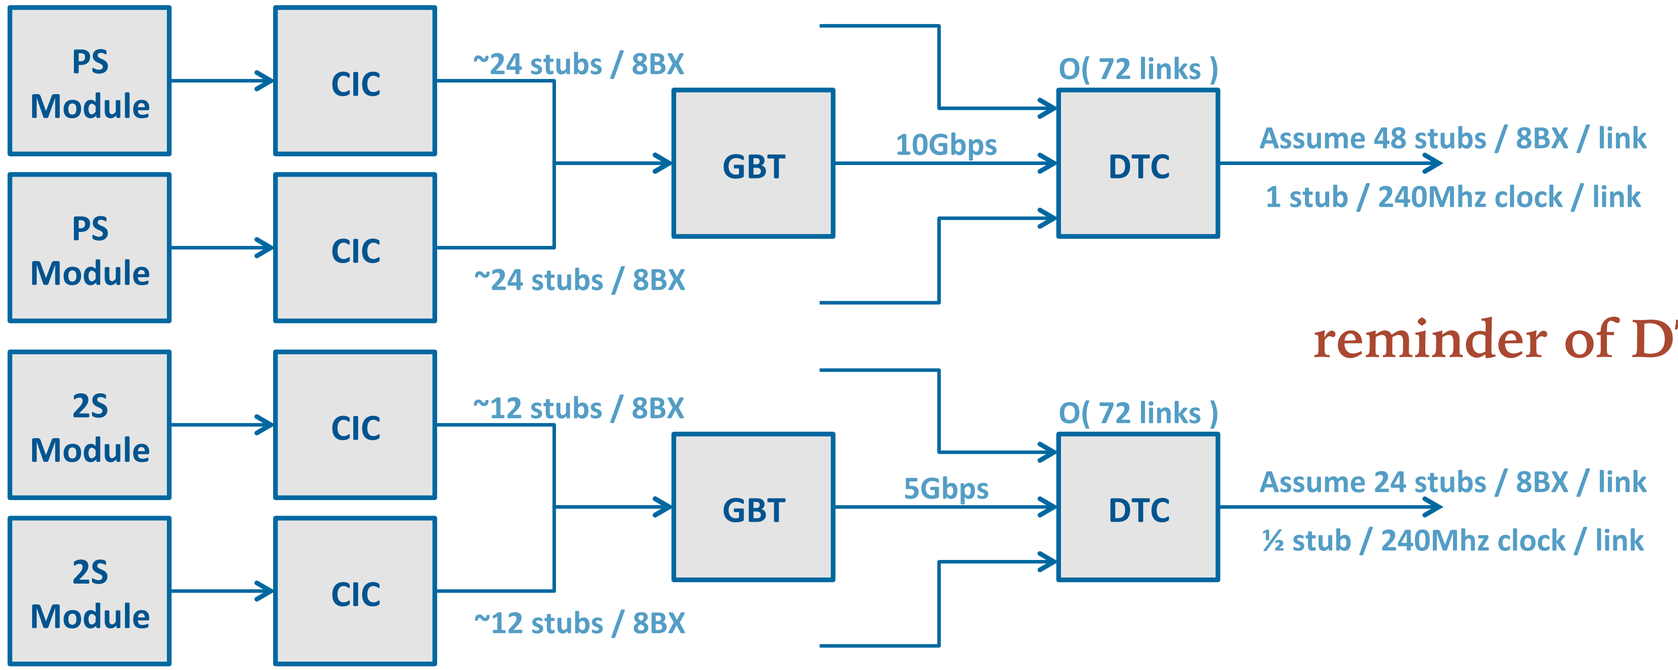


reminder of DTC input

# PROPOSED DTC - IMPLEMENTATION & LATENCY ESTIMATE

- We have proposed a DTC solution that provides us with our time multiplexed streams, but also avoids large fan-outs and fan-ins at all costs
- Objective is to place no requirements on the FPGA used, and to minimise the latency taken

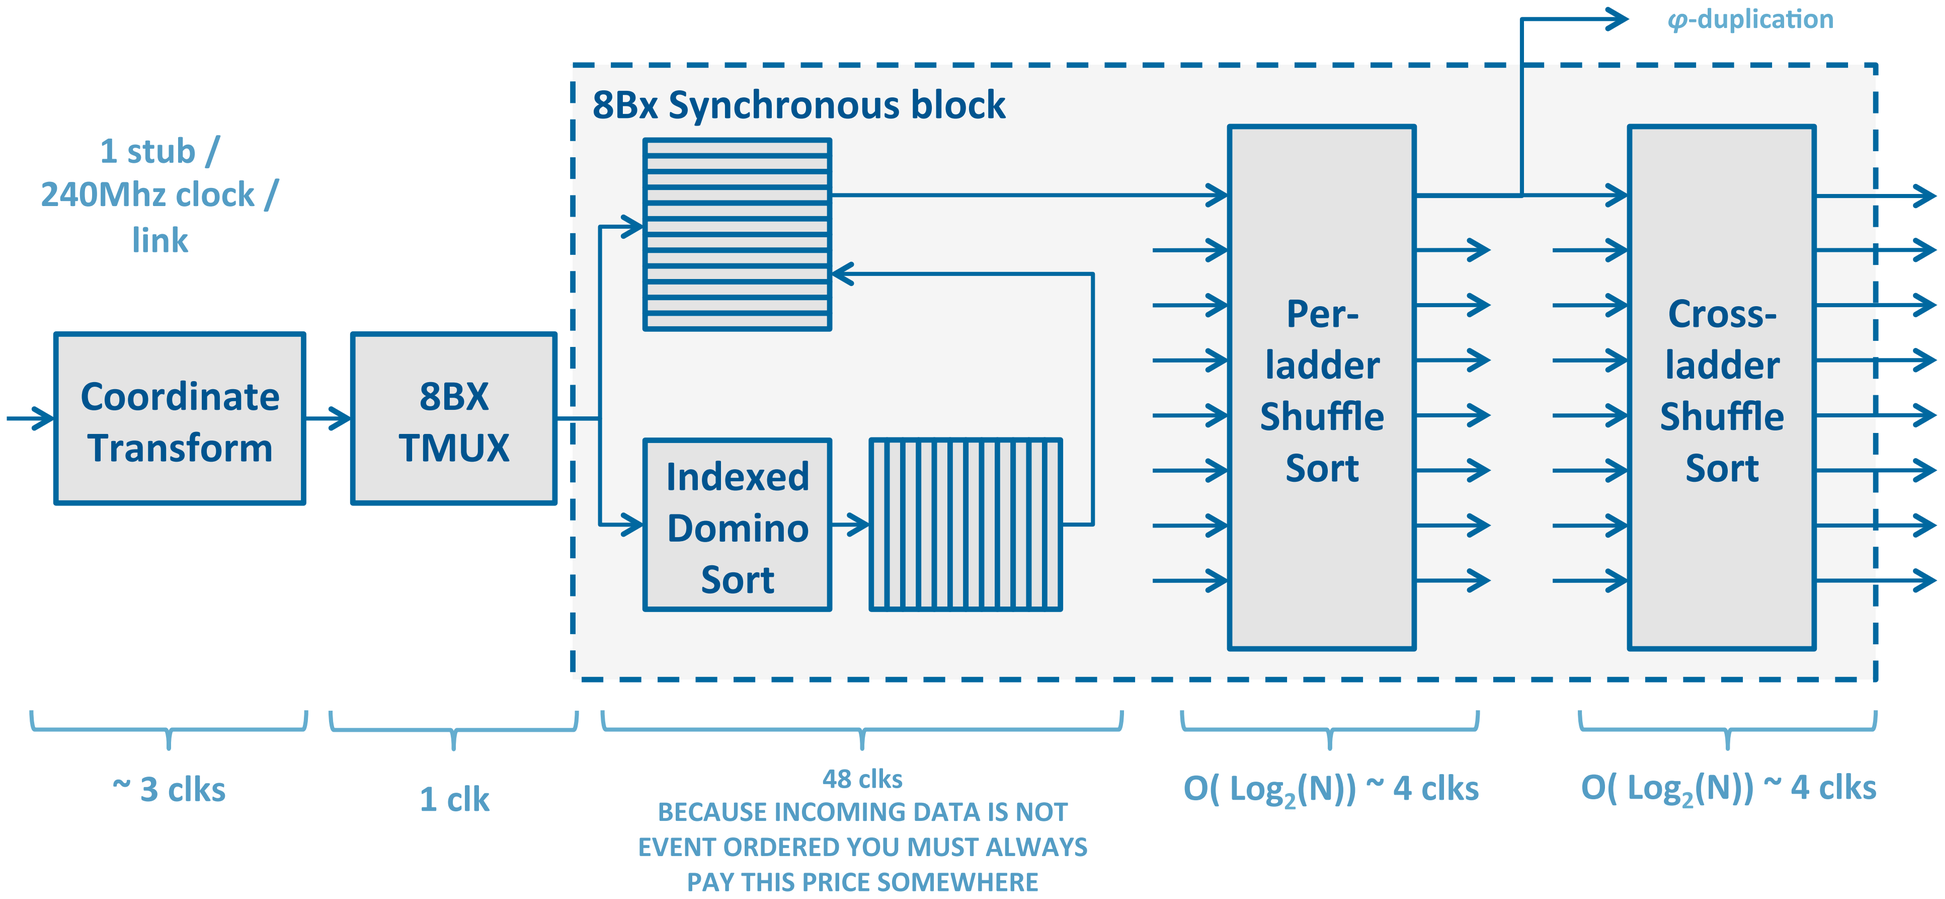


**DTC Latency estimate:**

**60 clocks**
**250 ns at 240 MHz**

Experience delivering a time-multiplexed calorimeter trigger and track-trigger demonstrator give us confidence that this is realistic

# PROPOSED DTC - IMPLEMENTATION & LATENCY ESTIMATE

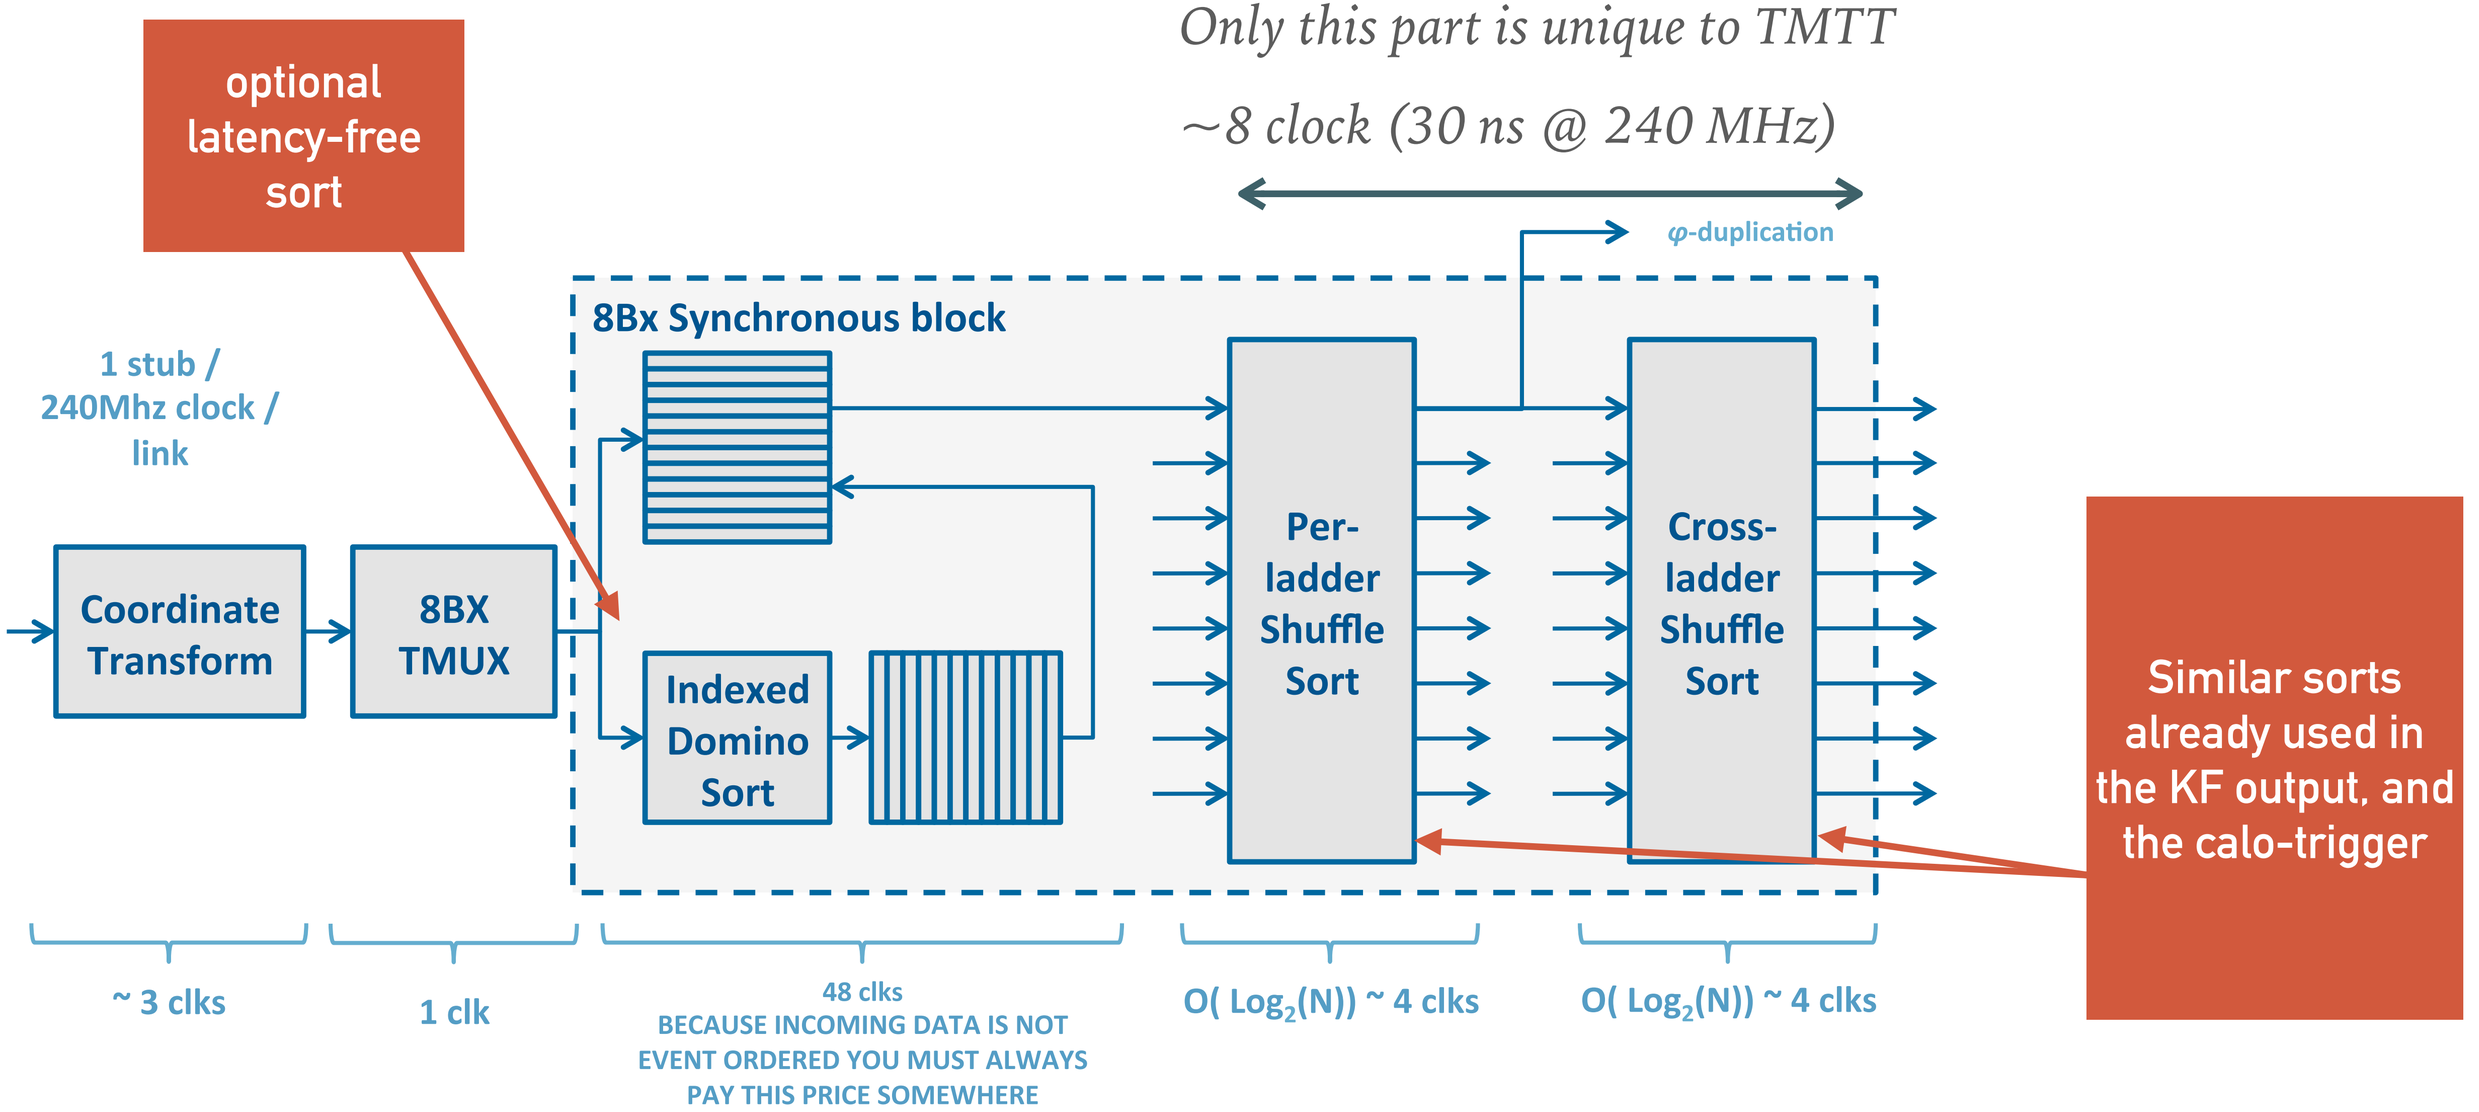

# DEMONSTRATOR LATENCY - MEASUREMENTS

- Latency of all parts of the demonstrator chain are **fixed**
  - **Independent of pileup or event**
- Latency measured for each block and set of links independently, and also of the total chain for validation

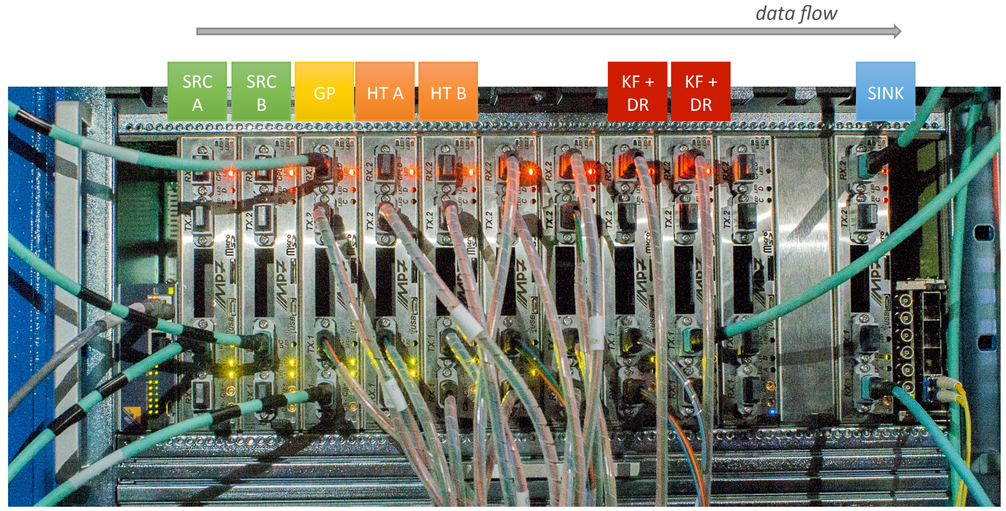


DTC estimate **250 ns**

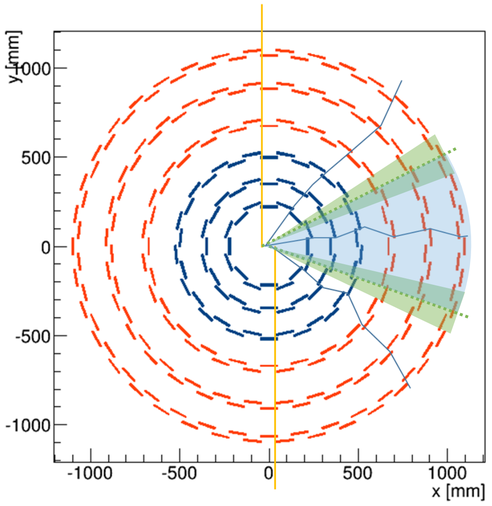


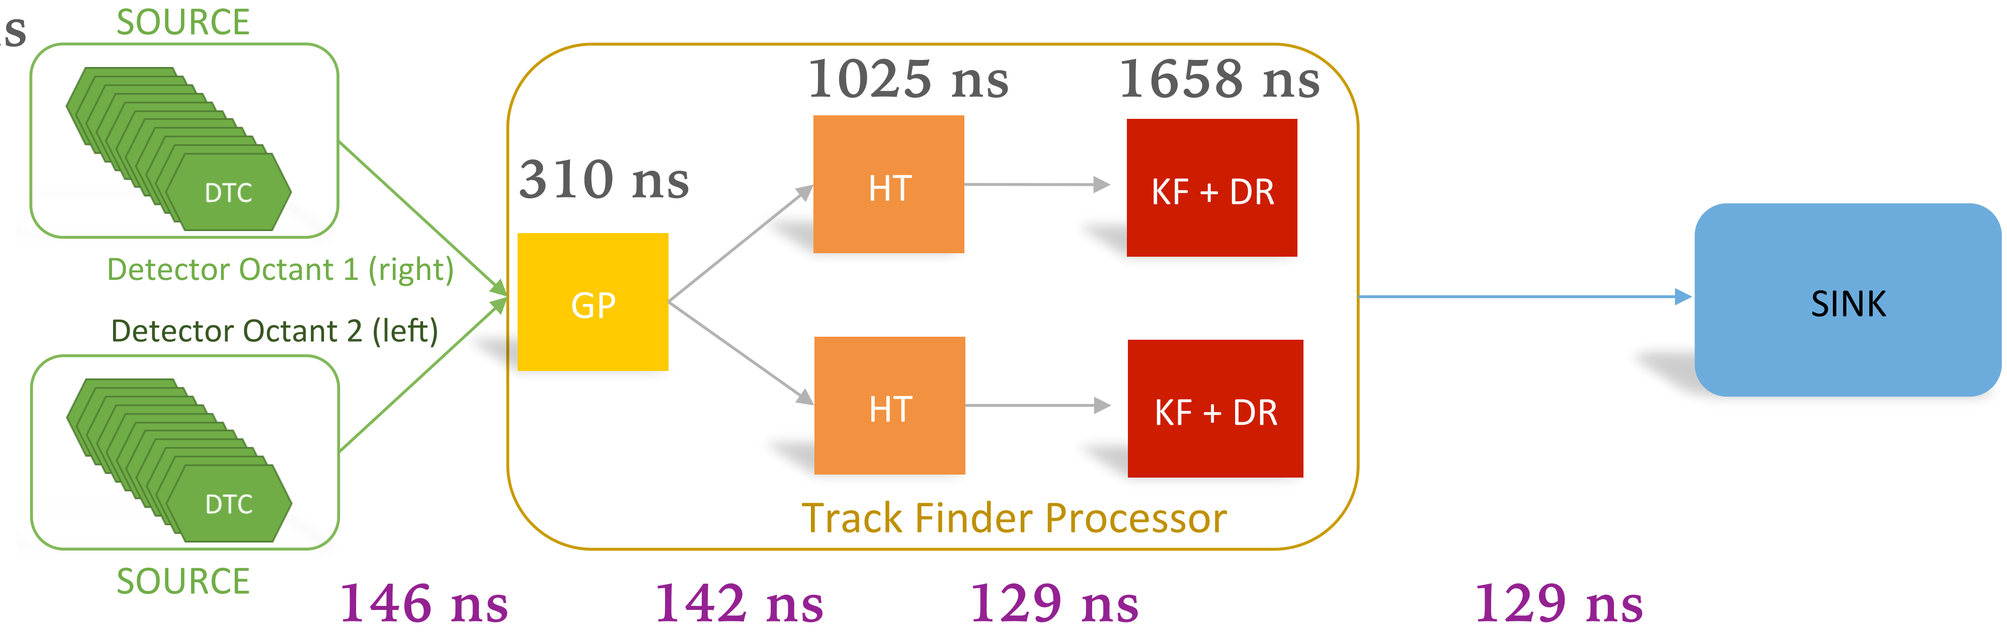

# DEMONSTRATOR LATENCY MEASUREMENTS

- Demonstrator target processing latency of 4 us has been achieved
- Latency has been tuned for worst case scenario (ttbar+200PU, flat tracker geometry)
- However, final system latency must also include the DTC, but fewer serdes & optics within the *Track Finding Processor*

| Demonstrator Chain | Latency (ns) |
|---|---|
| Serdes & optical length 1 | 143 |
| Geometric Processor | 310 |
| Serdes & optical length 2 | 144 |
| Hough Transform | 1025 |
| Serdes & optical length 3 | 129 |
| Kalman Filter + Duplicate removal | 1658 |
| Serdes & optical length 4 | 129 |
| **Total First out - First in** | **3538** |
| Last out - First out | 225 |
| **Total Last out - First in** | **3763** |

# DEMONSTRATOR LATENCY MEASUREMENTS

- Final system latency must also include the DTC, but fewer serdes & optics within the *Track Finding Processor*
- Have already explored ways we could reduce the latency further
  - All parts of system currently clocked at 240 MHz
  - Accumulation periods as we wait for all data to arrive in HT or KF could be reduced with faster link/smaller TMUX

| System Latency | Latency (ns) |
|---|---|
| DTC estimate | 250 |
| Serdes & optical length x 3 | 450 |
| Geometric Processor | 310 |
| Hough Transform | 1025 |
| Kalman Filter | 1620 |
| Duplicate removal | 38 |
| **Total First out - First in** | **3693** |
| Last out - First out | 225 |
| **Total Last out - First in** | **3918** |

# SUMMARY

- Have built a track-finding & fitting hardware demonstrator with currently available MicroTCA boards
- Capable of finding and fitting real physics tracks
  - over the entire 2π and $|\eta| < 2.4$ solid angle
  - one octant in φ at a time
- Have demonstrated
  - high efficiency and rate reduction in Monte-Carlo physics events
  - including TTbar + 200 PU
- With a fixed processing latency < 4.0 us

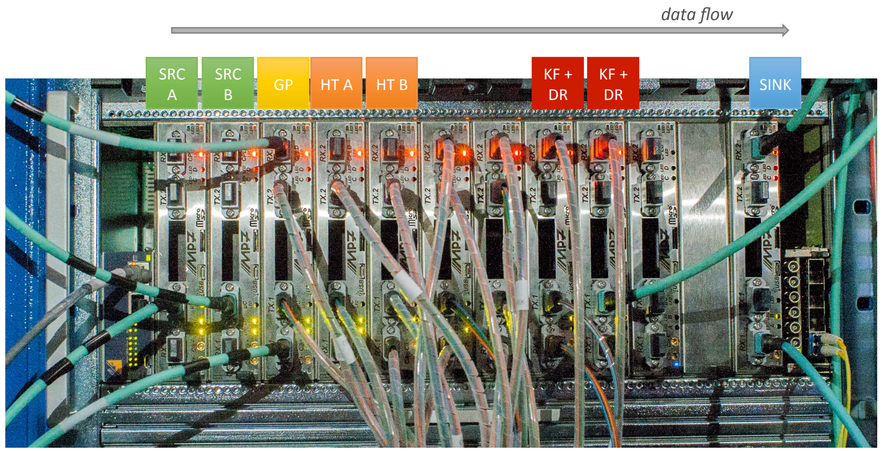

# BACKUP - DEMONSTRATOR RESULTS - RATE REDUCTION

➤ Hough Transform does the vast majority of the rate reduction

200 pileup

*~20,000 stubs in tracker*

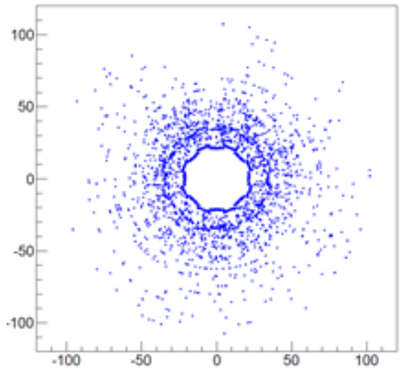

Hough Transform

*~300 track candidates*

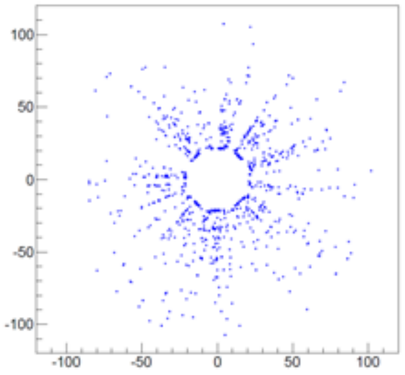

*On average 90 stubs enter HT*

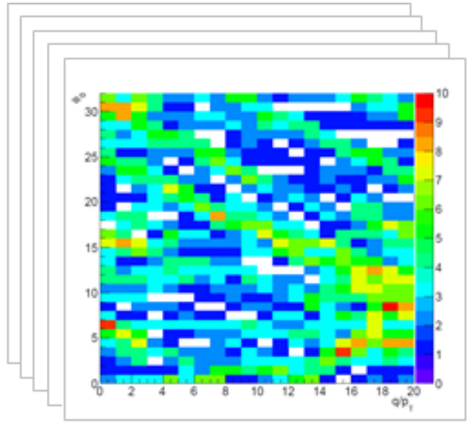

Hough Transform

*On average 6 leaves*

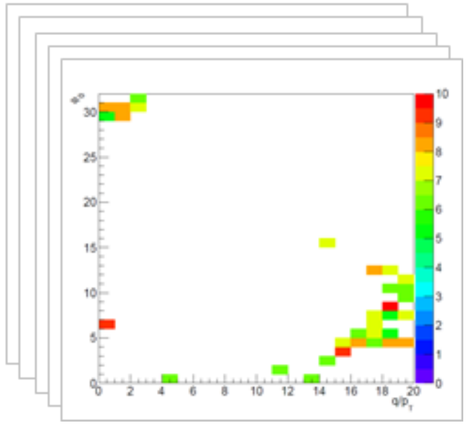

Jet pre bend filter