Readout firmware of the Vertex Locator

for LHCb Run 3 and beyond

Karol Hennessy

on behalf of LHCb
University of Liverpool / CERN
October 14, 2020

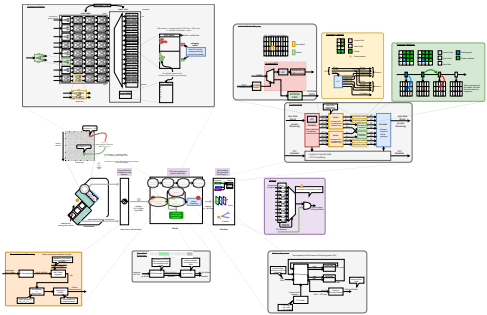

LHCb VELO

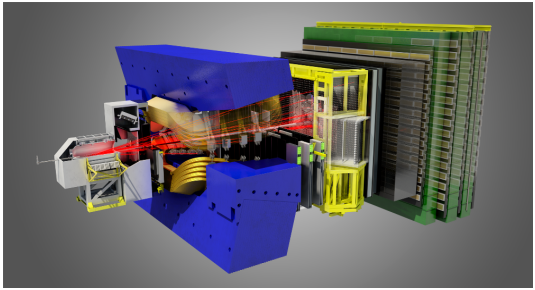

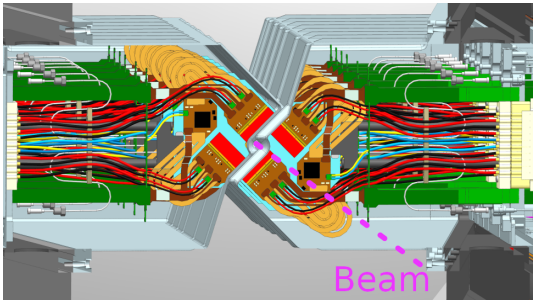

- LHCb - Flavour physics detector
- Excellent **vertexing** resolution and Particle ID
- LHCb has triggerless readout - full detector readout @ 40 MHz

- Vertex Locator (VELO)
- Silicon pixel modules around the LHC collision region
  - 50fb^{-1} integrated luminosity for LHC Runs 3 & 4
  - Very high radiation environment
  - In vacuum and under active cooling

LHCb VELO

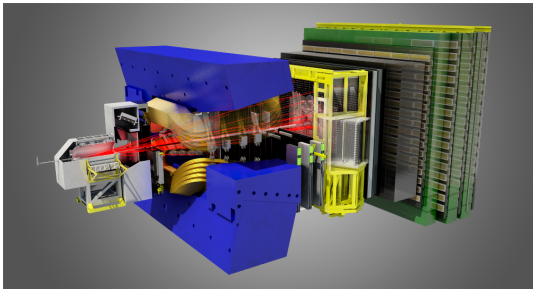

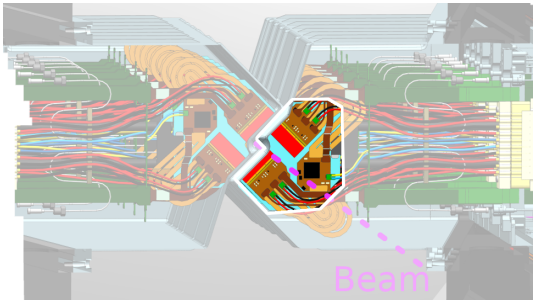

- LHCb - Flavour physics detector
- Excellent **vertexing** resolution and Particle ID
- LHCb has triggerless readout - full detector readout @ 40 MHz

- Vertex Locator (VELO)
- Silicon pixel modules around the LHC collision region
  - 50fb^{-1} integrated luminosity for LHC Runs 3 & 4
  - Very high radiation environment
  - In vacuum and under active cooling

VELO Module

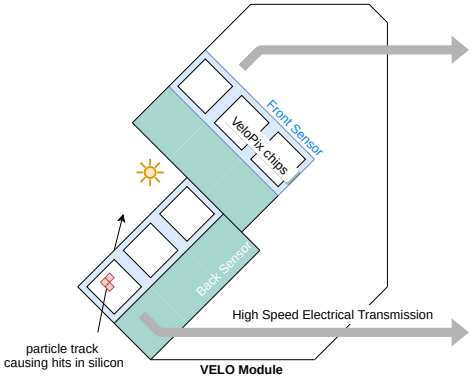

- Whole VELO = two halves of 26 modules
- Four sensors per module
  - 2 front
  - 2 back
- 3 VeloPix per sensor (i.e., 12 total)
- 20 high speed readout links
  - Chips closer to beam see more hits, and need more bandwidth

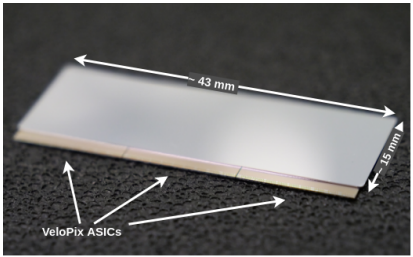

VELO Module

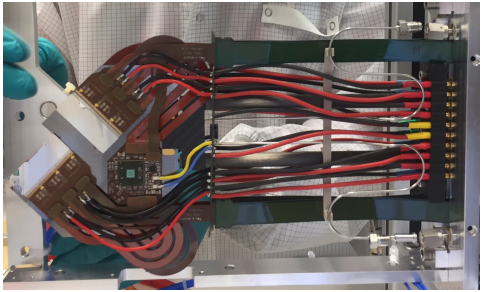

- Whole VELO = two halves of 26 modules
- Four sensors per module
  - 2 front
  - 2 back
- 3 VeloPix per sensor (i.e., 12 total)
- 20 high speed readout links
  - Chips closer to beam see more hits, and need more bandwidth

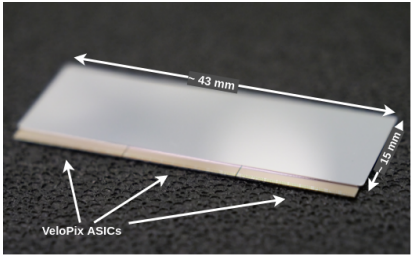

VeloPix ASIC

- Front-end ASIC driving the design of the VELO data acquisition system
- Part of the MediPix/TimePix family
- 130 nm CMOS technology
- 256×256 pixels of 55×55µm^2
- Clocked at 40 MHz
- Sends binary hit information (reducing bandwidth requirement)
  - full signal amplitude (ToT) available via slow readout for calibration

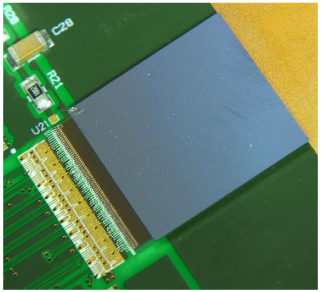

VELO Electronics and DAQ

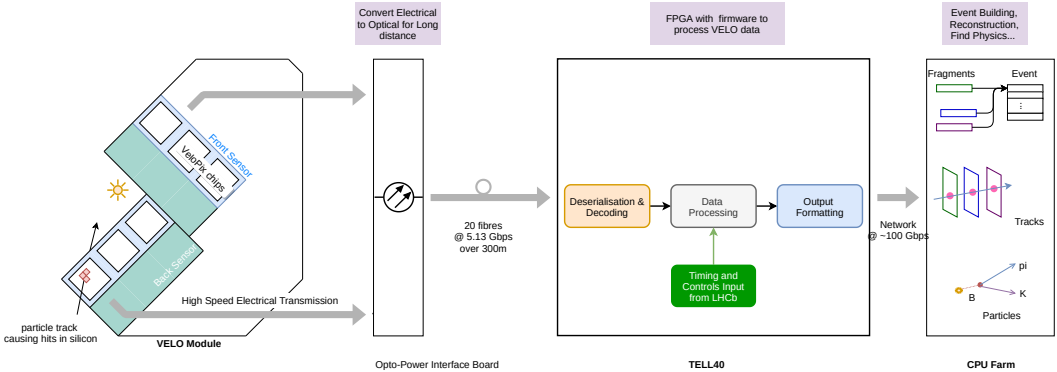

a slice of the VELO readout system

- See Flavio's talk on Friday for a fuller description of the LHCb DAQ

Why an FPGA?

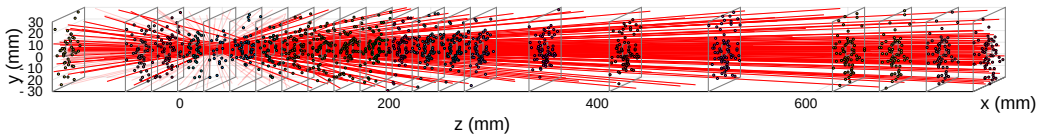

- *Lots of data!*
- VeloPix is optimised for high speed readout

Peak hit rate	900 Mhits/s/ASIC
Max data rate	19.2 Gb/s
Total VELO	2.85 Tb/s

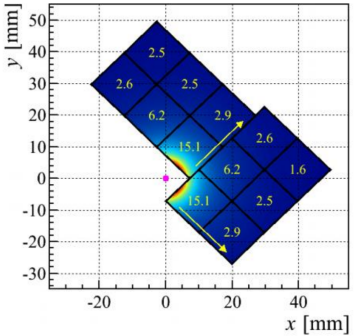

Data rate [Gbit/s] for hottest module.

Readout Board - PCIe40/TELL40

- **Single control and readout board for the entire experiment**
- Can be used for Timing, Slow Control, DAQ or all
- Common hardware, shared firmware components
- PCIe Gen3 x16
- **Intel Arria10 FPGA** (10AX115S4F45E3SG)
- 1 TELL40 = 1 VELO module

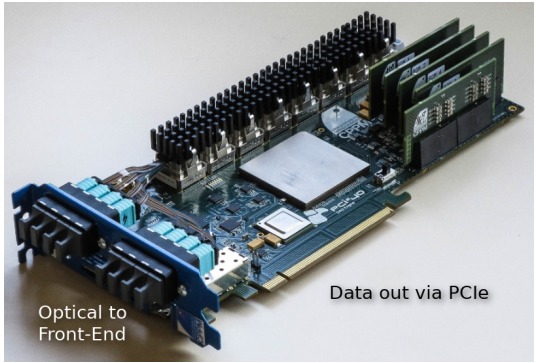

- up to 48 bi-directional links @ $\sim$5 Gb/s
- Output bandwidth 100 Gb/s (measured).

VELO Electronics and DAQ

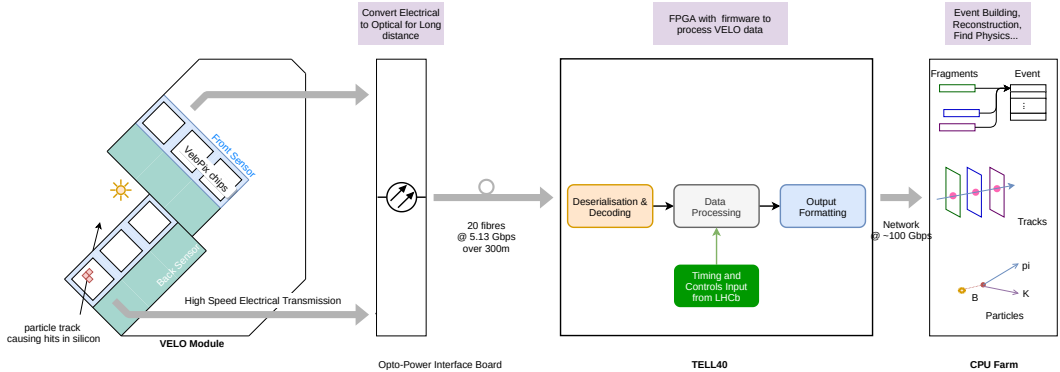

So what does the VELO Firmware have to do?

What does VeloPix produce?

- **Time unordered data**

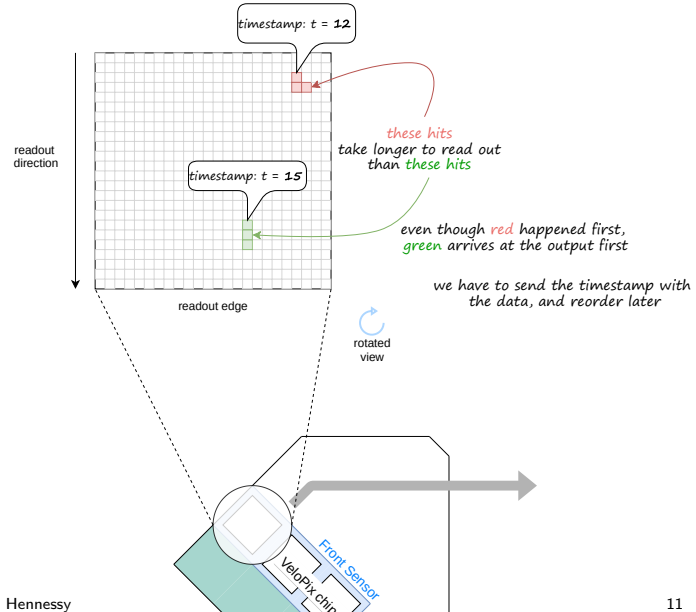

What does VeloPix produce?

- Time unordered data
- **Custom transmission Protocol**

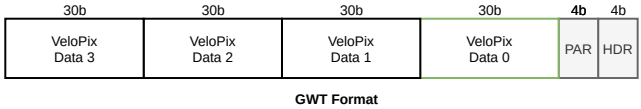

GWT Format

- Custom serializer - **Gigabit Wireline Transmitter (GWT)**
  - Chosen for low power - 60 mW
  - 5.12 Gb/s line rate (slightly higher than 4.8 Gb/s of G**B**T)
- GWT protocol
  - scrambled data (30 bit multiplicative)
  - parity check, no error recovery
  - low tolerance for header errors

What does VeloPix produce?

- Time unordered data
- Custom transmission Protocol
- **SuperPixels**

- Pixel data is aggregated into groups of 2×4 called **SuperPixels**
  - 30% reduction in data size
- Timestamp stored in SuperPixel data packet

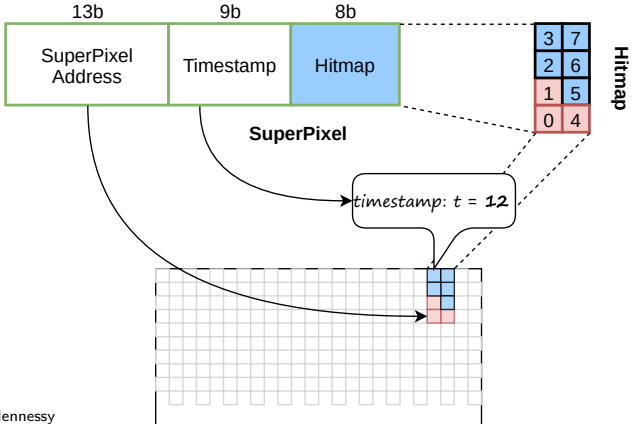

Readout Firmware

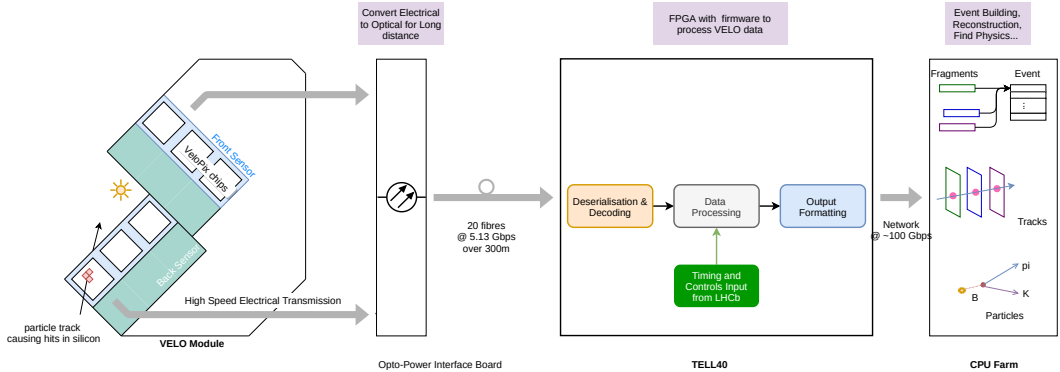

What's in the TELL40 Firmware?

Readout Firmware

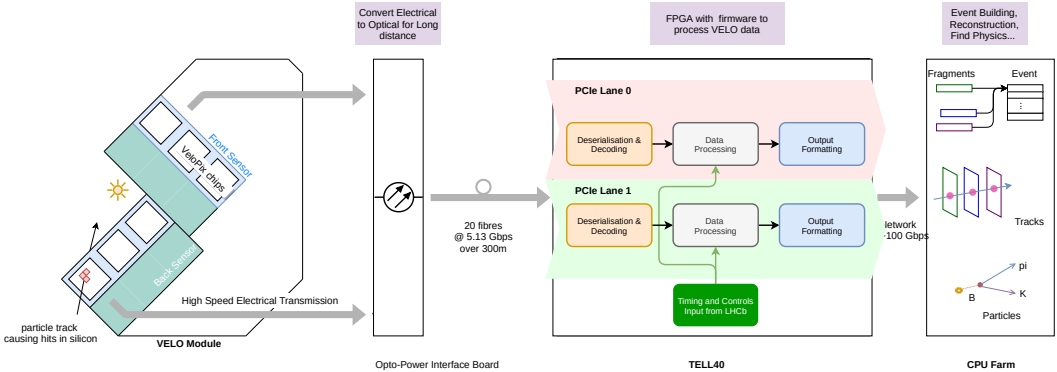

- Actually, it's two parallel streams for PCIe bandwidth optimisation
- But it's simpler to describe just one

Handling VeloPix data

Going back to our list:

1. Custom transmission Protocol
2. Time unordered data SuperPixels
3. *is a generic component and won't be discussed here*

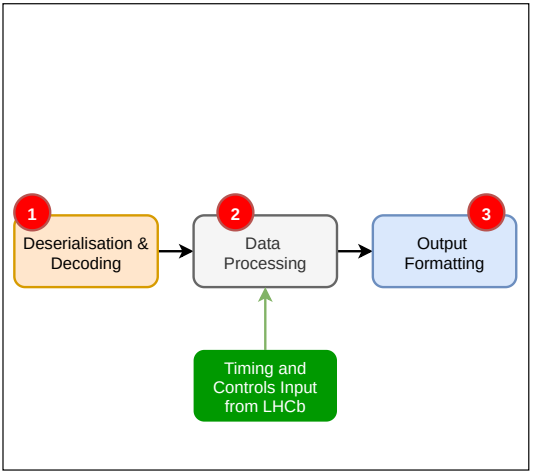

Handling VeloPix data - Deserialisation & Decoding

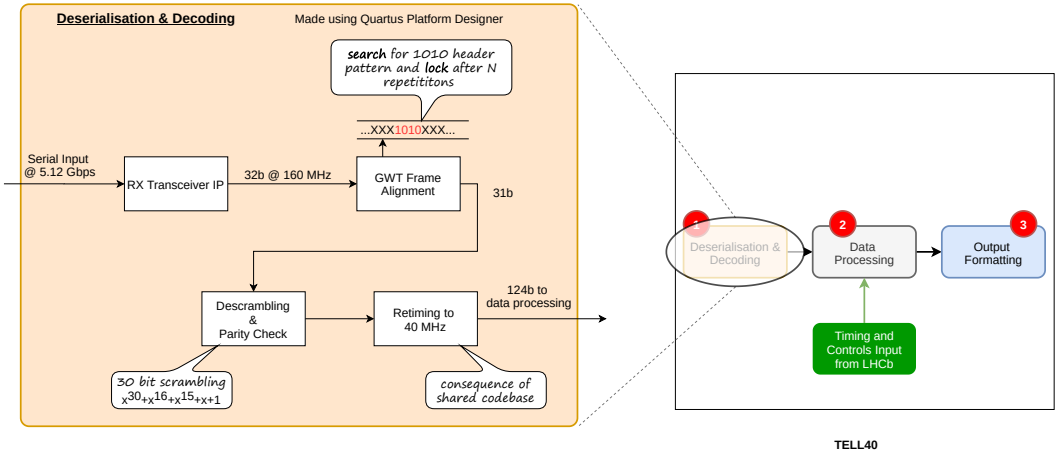

Handling VeloPix data

Going back to our list:

1. Custom transmission Protocol
2. Time unordered data SuperPixels
3. *is a generic component and won't be discussed here*

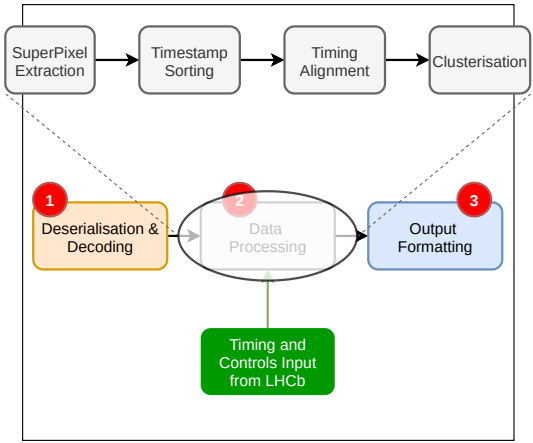

Handling VeloPix data - SuperPixel Extraction

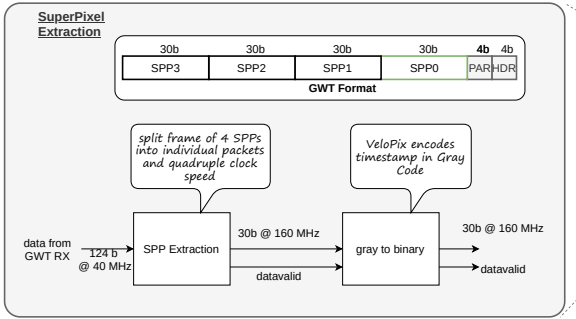

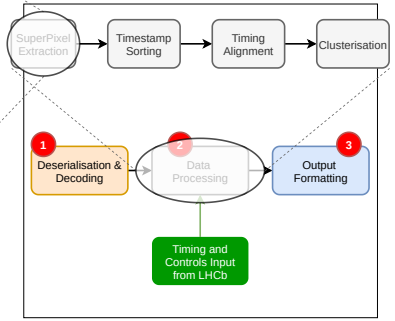

Handling VeloPix data - Time Reordering

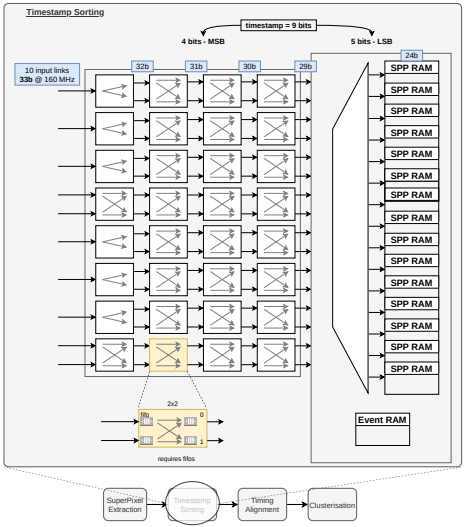

- Timestamps are sorted 1 bit at a time in several layers
  - First column is MSB...
- Fifos are needed to avoid collisions
- Data are stored in RAMs at the end of the routing
- The whole reordering consumes a large amount of the FPGAs memory

Handling VeloPix data - Time Reordering

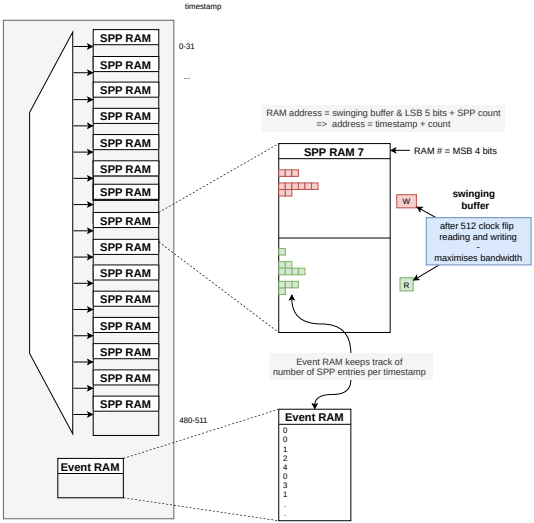

- After the routing, the RAM address is equivalent to the timestamp
- SuperPixel Packets are stored for a maximum latency of 512 clock cycles
- A swinging buffer is used to maximise bandwidth

Handling VeloPix data - Time Alignment

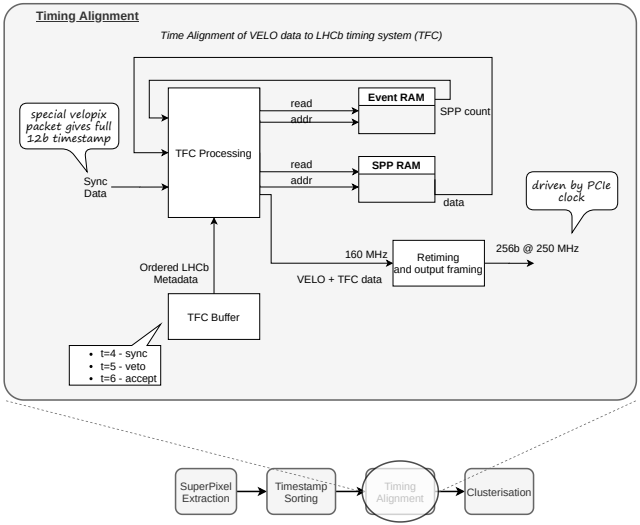

- After reordering, data must align to the rest of LHCb
- **Timing and Fast Control** (TFC) system provides LHCb timing metadata

Handling VeloPix data - Clusterisation

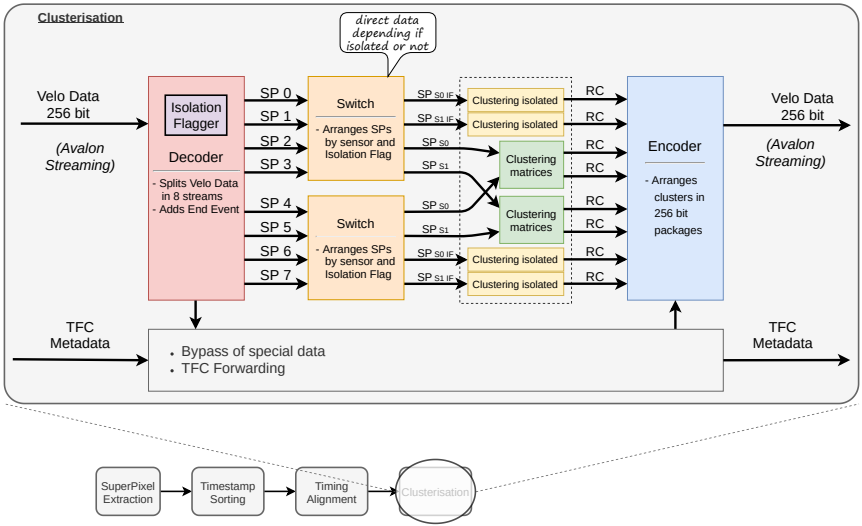

Handling VeloPix data - Clusterisation

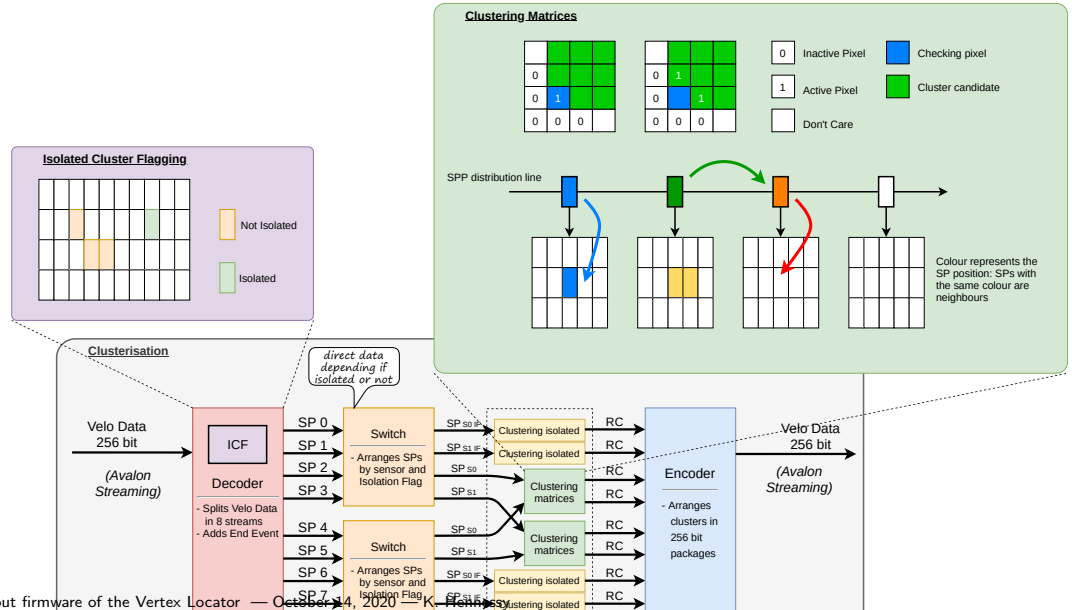

Challenges

Deserialisation & Decoding - Congestion

v1

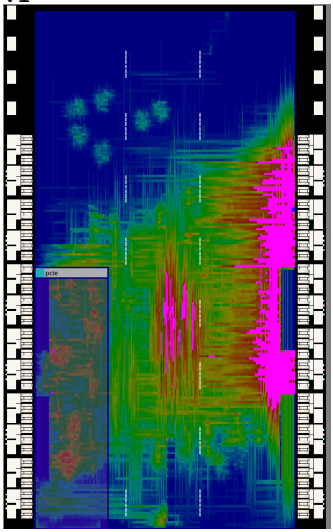

v2 Improved frame aligner

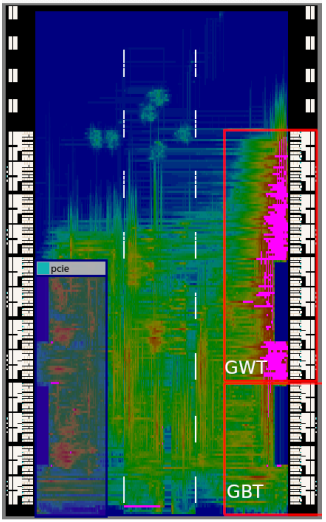

v3 Improved bit slipping

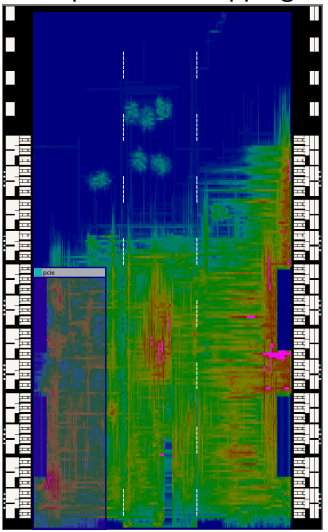

Full Data Processing

Now adding the full Time Reordering and Alignment

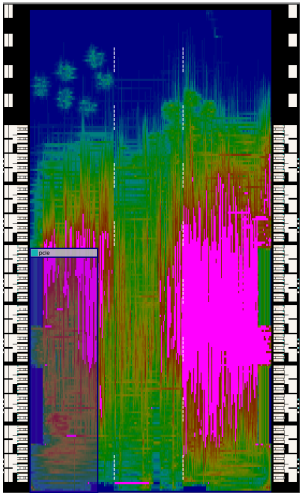

Slow 900mV 100C Model Setup Summary

<<Filter>>

	Clock	Slack	End Point TNS
1	altera_reserved_tck	-42.950	-42.950
2	lli_inst|\multiLink_gen_loop:2:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[2]|rx_pma_clk	-8.931	-68.411
3	lli_inst|\multiLink_gen_loop:2:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[3]|rx_pma_clk	-7.072	-28.643
4	lli_inst|\multiLink_gen_loop:2:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[5]|rx_pma_clk	-6.408	-25.894
5	TELL40_1|\GEN_FULL_DP:data_proc|\pll_dp_a10_gen:inst_Data_Processing_clock|iopll_0|outclk1	-5.173	-72643.573
6	lli_inst|\multiLink_gen_loop:2:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[0]|rx_pma_clk	-4.718	-17.192
7	TELL40_0|\GEN_FULL_DP:data_proc|\pll_dp_a10_gen:inst_Data_Processing_clock|iopll_0|outclk1	-4.530	-61712.754
8	C100_osc	-2.237	-57.151
9	lli_inst|\multiLink_gen_loop:4:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[3]|rx_pma_clk	-1.559	-4.695
10	lli_inst|\multiLink_gen_loop:3:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[0]|rx_pma_clk	-1.321	-1.359
11	lli_inst|\multiLink_gen_loop:2:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[1]|rx_pma_clk	-1.305	-17.491
12	lli_inst|\multiLink_gen_loop:3:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[1]|rx_pma_clk	-1.166	-3.869
13	lli_inst|\multiLink_gen_loop:3:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[3]|rx_pma_clk	-0.946	-1.879
14	lli_inst|\multiLink_gen_loop:2:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[4]|rx_pma_clk	-0.889	-1.923
15	pcie_top|pcie_1|qsys_pcie|pcie|coreclkout	-0.598	-14.636
16	lli_inst|\multiLink_gen_loop:4:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[4]|rx_pma_clk	-0.566	-5.660
17	pcie_top|pcie_0|qsys_pcie|pcie|coreclkout	-0.399	-9.267
18	pcie_top|pcie_0|pcie_pll|iopll_0|outclk_220	0.029	0.000
19	lli_inst|TFC_XCVR1_inst|xcvr_native_8b10b_deterministic_latency_cpri|rx_clkout	0.215	0.000
20	lli_inst|\multiLink_gen_loop:5:multiLink_gen_gw..._native_a10_0|g_xcvr_native_insts[0]|rx_pma_clk	0.436	0.000

Timing closure becomes tricky

Resource Estimate

- Very preliminary estimate for now (sum of individual compilations - not the output of a complete build)

	Logic (ALMs)	M20K RAMs
Timestamp Reordering	39	72
Clustering	31	11
Total	70	83

- Looks like it will fit, but **congestion** and **timing closure** are the major challenges ahead

A working firmware...

- Full data processing is not complete today
- A bypass is used for production and testing
  - Sorting/processing is done on CPU
  - Rate limited
- Can also be used to check data processing (send same data to both and check)

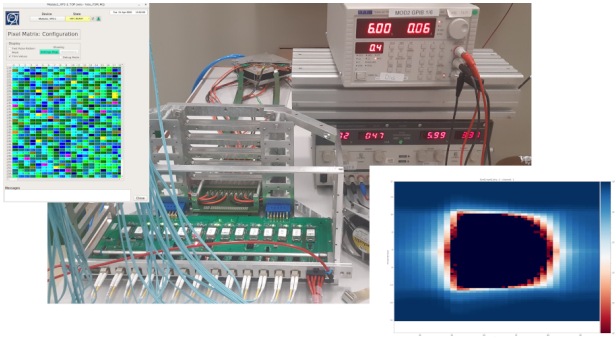

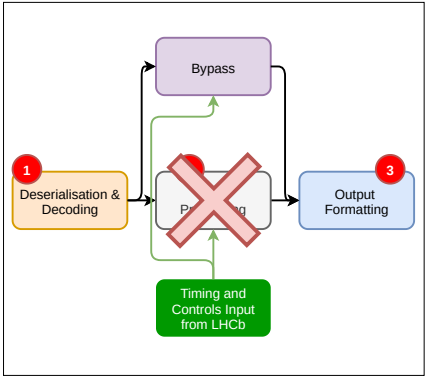

Tools/Organisation

- Typical combination of Questasim and Quartus
- LHCb employs gitlab pipelines for checking new releases
  - Sim-checker injects files into firmware and verifies on output
  - Can add additional testbenches and cross-checks à la “nightlies”
  - Strict versioning and tracking
- VELO makes stable releases for production testing

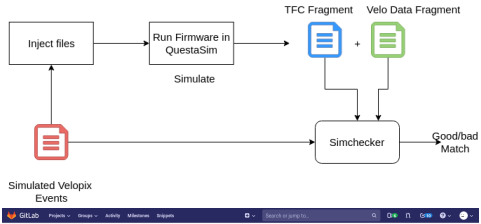

...and beyond

Timeline

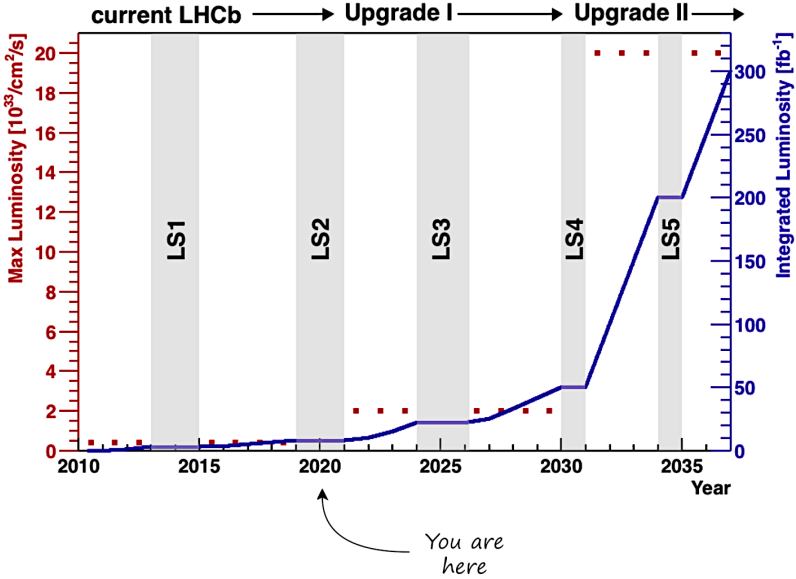
current LHCb
Upgrade I
Upgrade II
Max Luminosity [10^33/cm^2/s]
Integrated Luminosity [fb^-1]
LS1
LS2
LS3
LS4
LS5
Year
You are here

VELO U2

- HL-LHC (2028) will provide 7.5× luminosity
- Meaning 7.5× tracks/hits...
- Meaning we need a new VELO to go from this

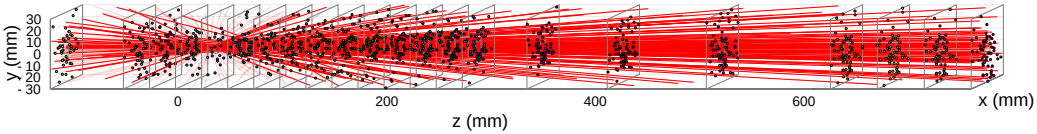

- to this

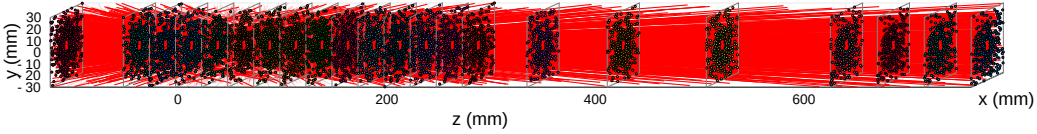

...and beyond

- Add extra timing precision (necessary for vertex/tracking)
- Bandwidth increase $O(10)$

	Arria10	Agilex*	Factor Increase
Process	20 nm	10 nm	~ 2
Logic Elements (k)	1150	2692	2.3
M20k Memory (Mb)	53	259	4.9
DSP	1518	17056*	~ 16

Table: Comparison of FPGA resources for VELO U1b and a candidate for U2.

- Next gen FPGAs not quite scaling with the needs of the experiment!
- What can we do with all these DSPs?

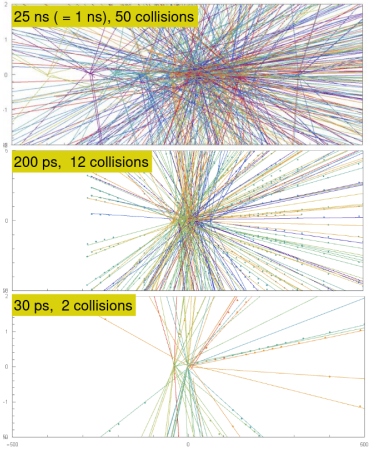

Concluding Remarks

- LHCb VELO firmware on track to process VeloPix data
- Validating Time Reordering and Clustering
- Several challenges in terms of FPGA resources and timing closure
  - Confident we can solve these
  - We welcome any clever suggestions/tips
- Learning techniques to optimise the next generation of the experiment
- Need to adapt to the changing landscape of heterogenous computing

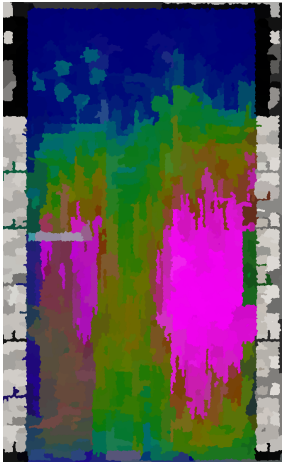

backup

BXID Router

- Time-ordering SuperPixel data
  - 9-bit **router** sorts data 1 bit at a time
  - Extensive simulation required - both to maximise speed (>160 MHz) and minimise FPGA resource usage
  - Latency limit < 512 clock cycles

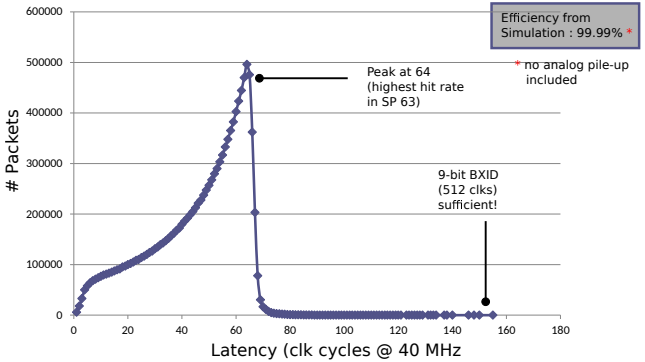

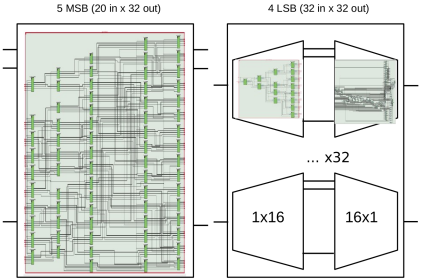

Special VeloPix

A lot of non-standard DAQ elements...

- VeloPix has **NO SCA**
  - SLVS communication component required
  - extra SOL40 firmware
  - extra SOL40 software
- VeloPix **does NOT use GBT** for DAQ
  - uses GWT
  - different frequency - 5.12 Gbps
  - special VELO LLI firmware component
  - special firmware decoding, clocking...
  - special VELO LLI software

- VeloPix sends data **"unsynchronised"**
  - Firmware re-aligns data by BX ID
  - Cannot filter events pre-alignment
  - Special dataflow monitoring needed
- Big effort from Online, Annecy, Marseille to help integrate into the standard firmware and software. Must remain vigilent and ensure "special cases" are tested as standard.

Bypass detail

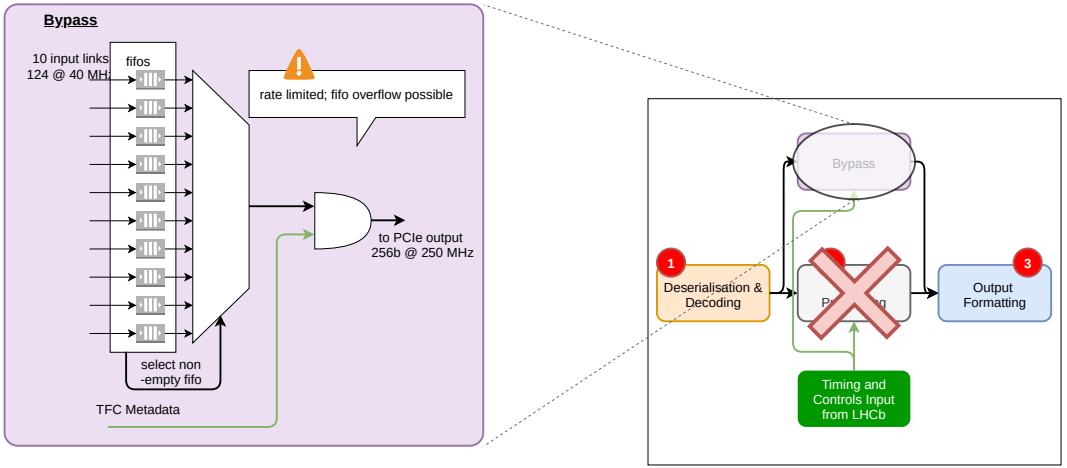

Isolated Cluster Flagging

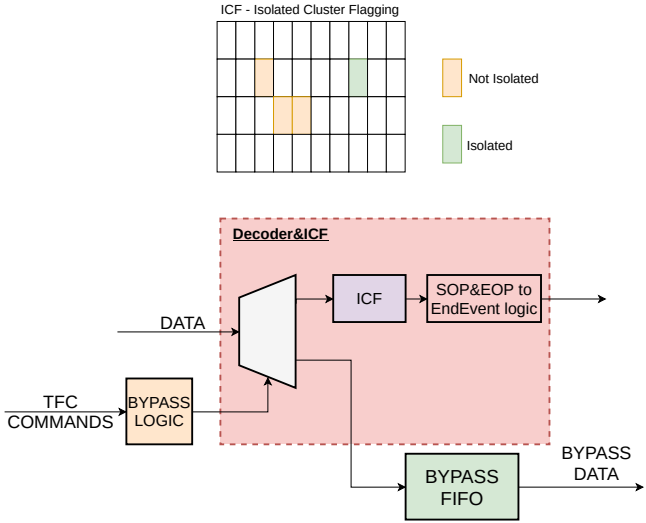

Isolated Clustering

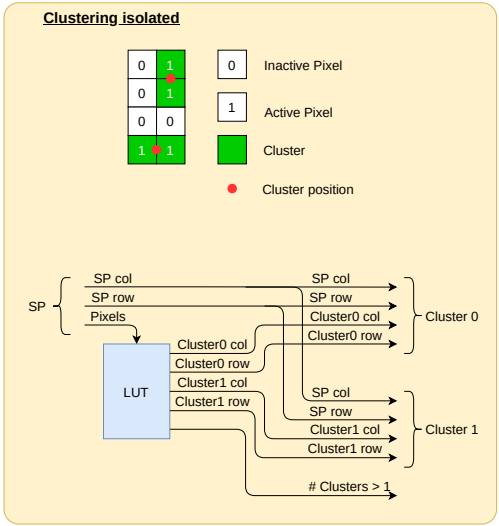

Clustering Matrices