Reconstructing jets in the Phase-2 upgrade of the CMS Level-1 Trigger with a seeded cone algorithm

Sioni Summers for the CMS Collaboration

8th May 2023

Introduction

- CMS Phase-2 Level-1 Trigger will have access to rich data from new detectors
  - Tracking for charged particles above 2 GeV
  - High Granularity Calorimeter for $|\eta| > 1.5$
  - Increased precision from barrel calorimeter to L1T, various upgrades to muon systems
- The offline Particle Flow and PUPPI algorithms have been ported to L1T FPGAs to make best use of the new data
  - Linking hits from subdetectors to make particle hypotheses (e.g. track & hadronic cluster → charged hadron)
  - Pileup Per Particle Identification (PUPPI) to mitigate contribution from 200 pileup vertices using particles' local environment
  - Executed in system Correlator Layer 1
  - Next talk in session for more details
- Correlator Layer 2 reconstructs higher-level quantities from PUPPI candidates
  - Jets, taus, energy sums (MET), electron/photon cleaning
- We present an algorithm for reconstructing jets in the Correlator Layer 2 FPGAs within a latency less than 1 μs

CMS Phase 2 Level 1 Trigger

- System will comprise 100s of FPGA boards reconstructing and reducing data in stages
- Correlator Layer 2 will use APx (left) and Serenity (right) ATCA boards with Xilinx Ultrascale+ FPGAs and multi-Tb/s data IO over optical fibres

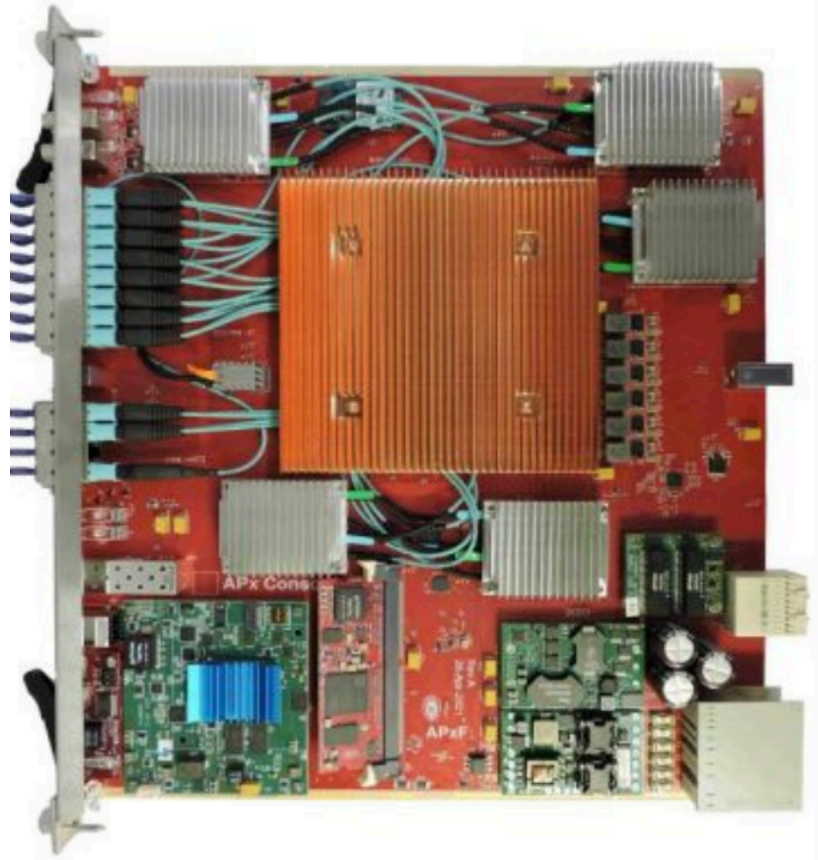

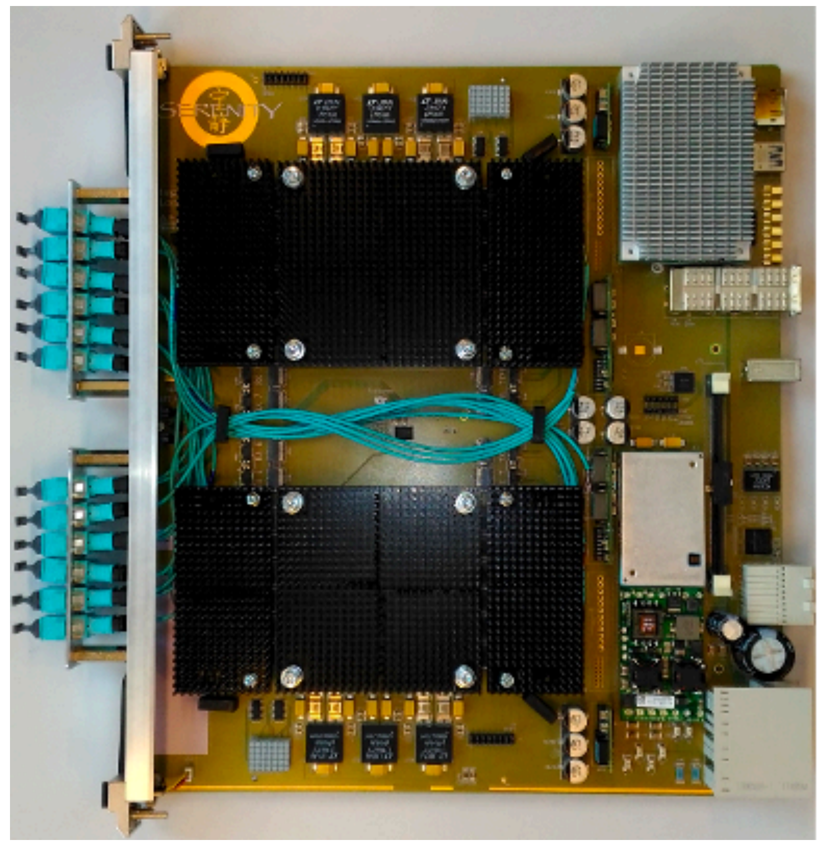

t = 0

Each small box is one Xilinx Ultrascale+ FPGA

t ≲ 12.5 μs

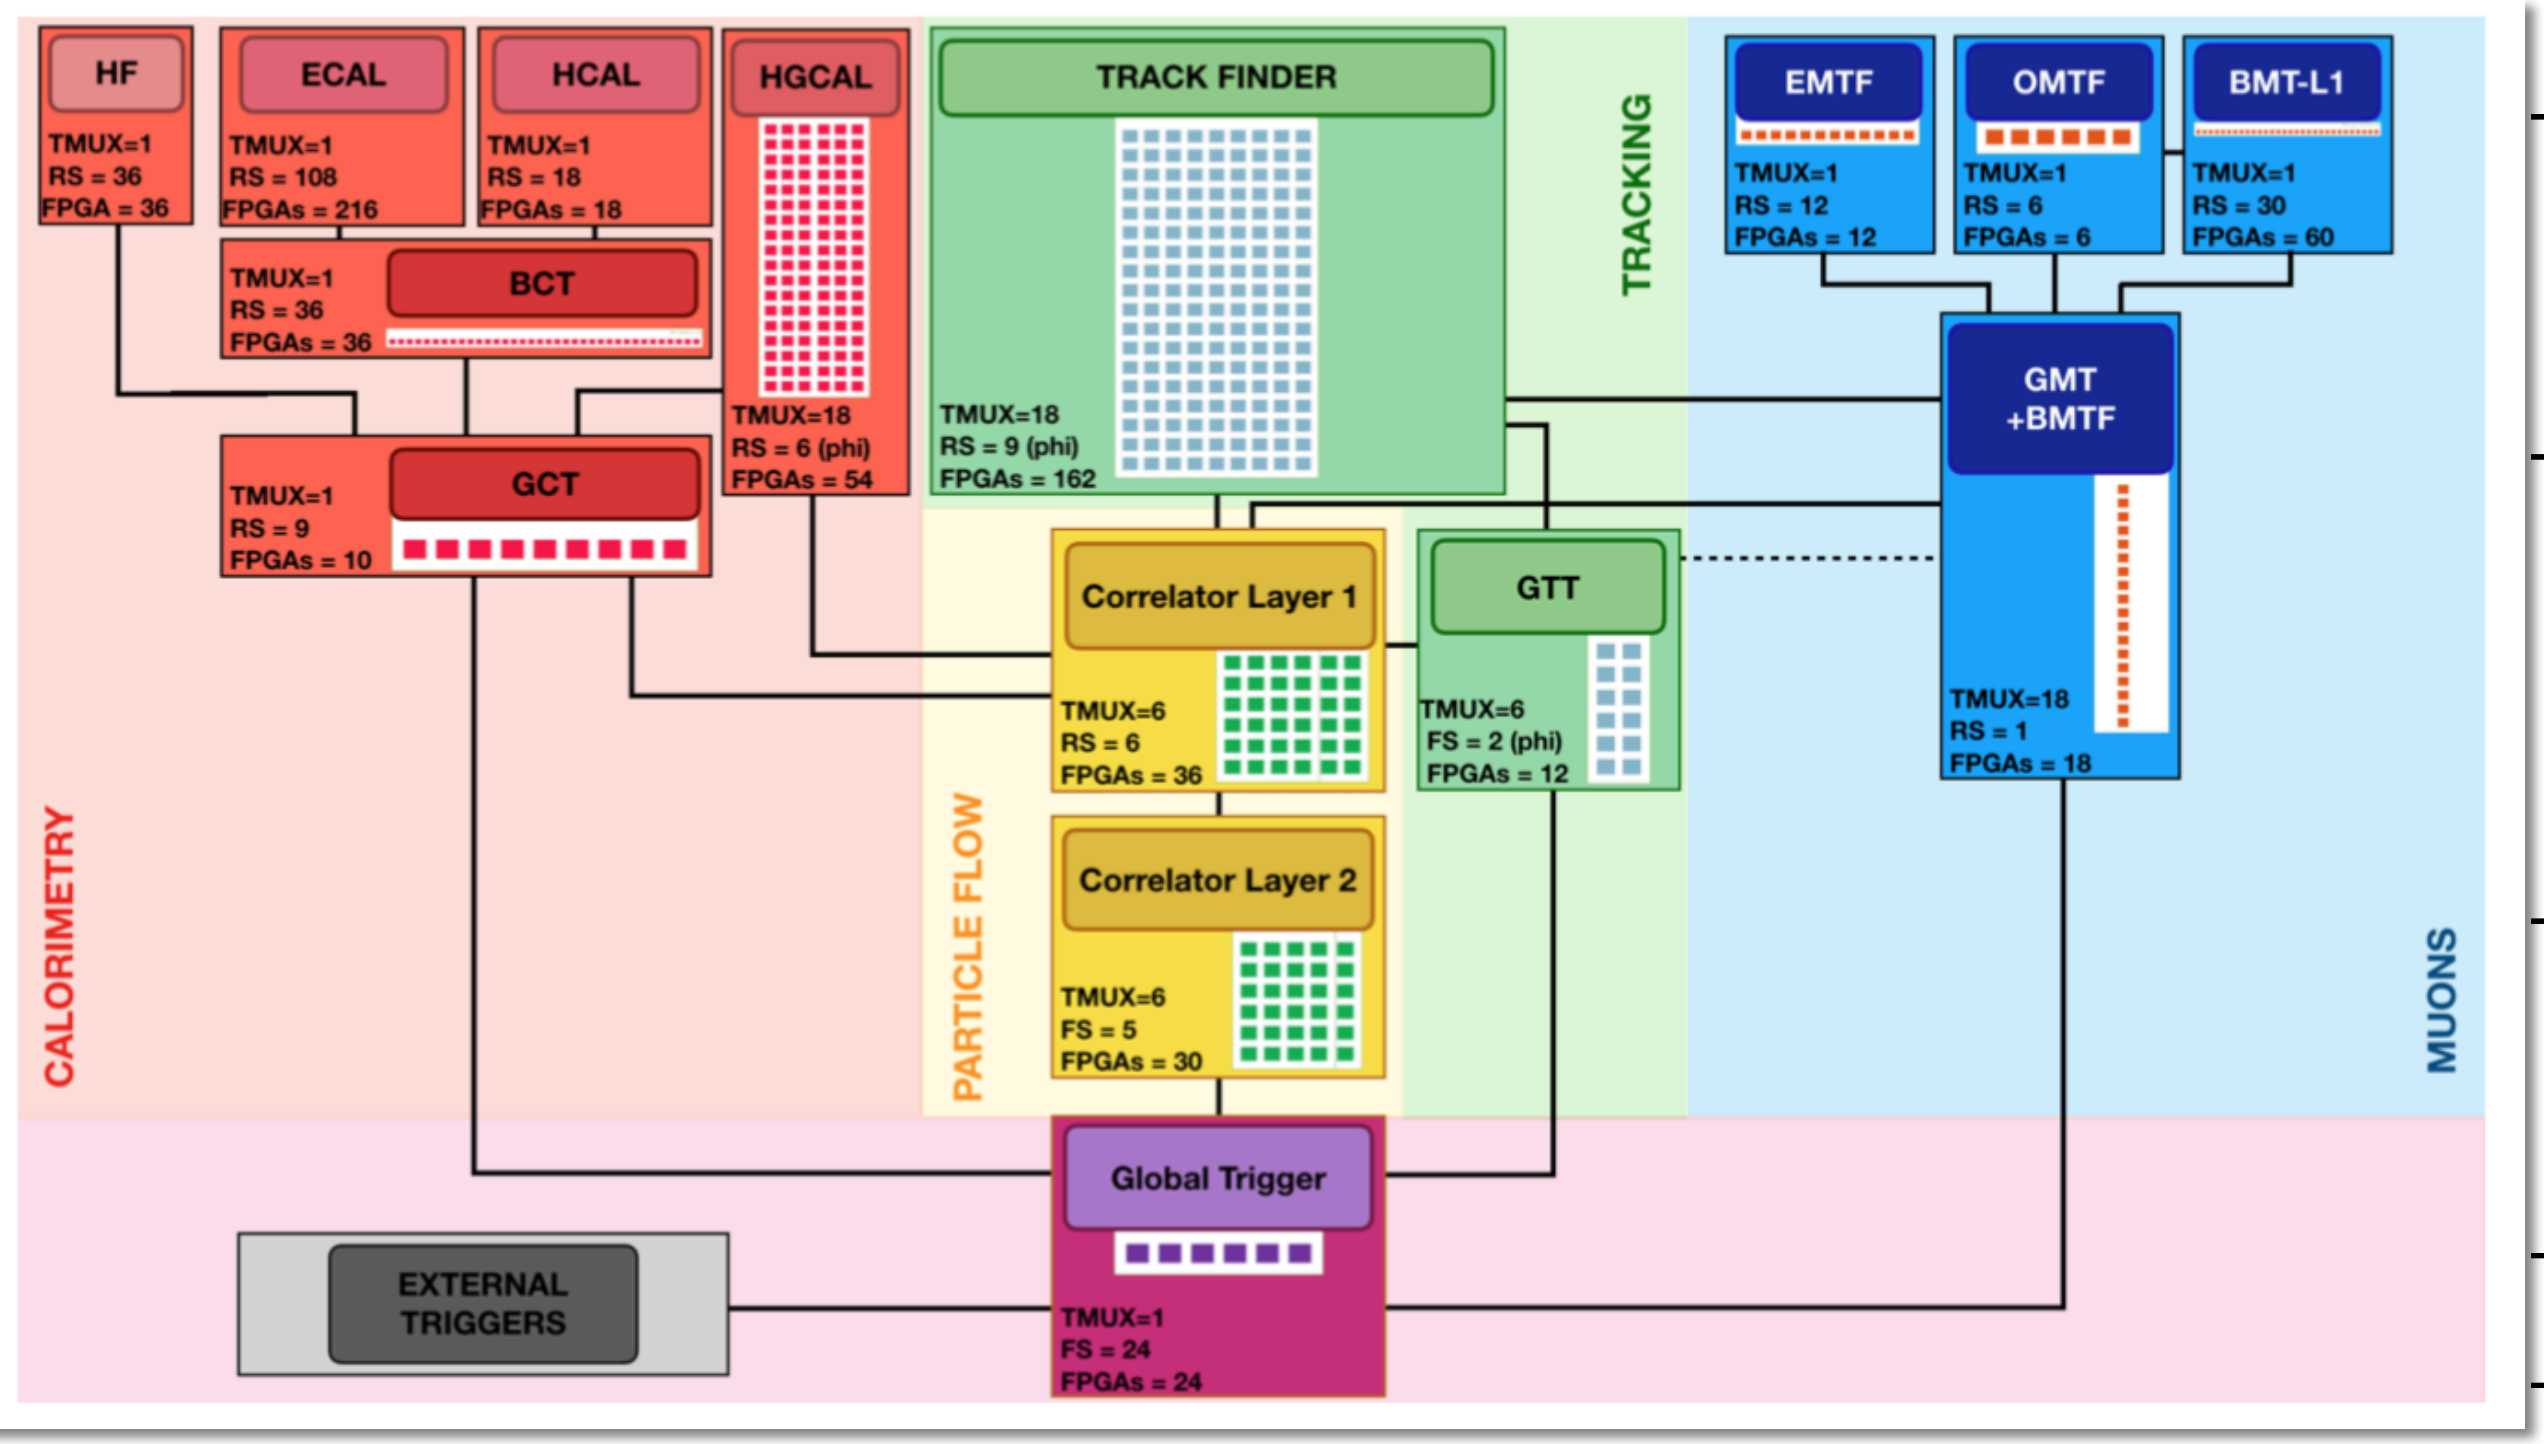

Detector hits

Clusters & Tracks

Particles

Event Categorisation

1 bit: keep / discard

CMS Phase 2 Level 1 Trigger

- System will comprise 100s of FPGA boards reconstructing and reducing data in stages
- Correlator Layer 2 will use Serenity and APx ATCA boards with Xilinx Ultrascale+ FPGAs and multi-Tb/s data IO over optical fibres

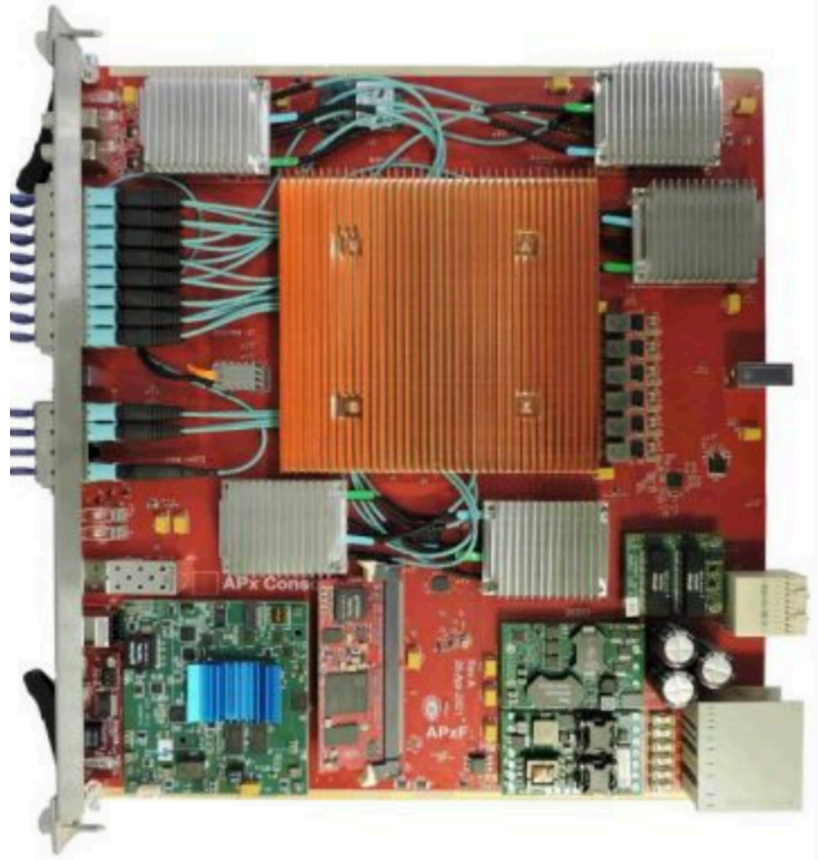

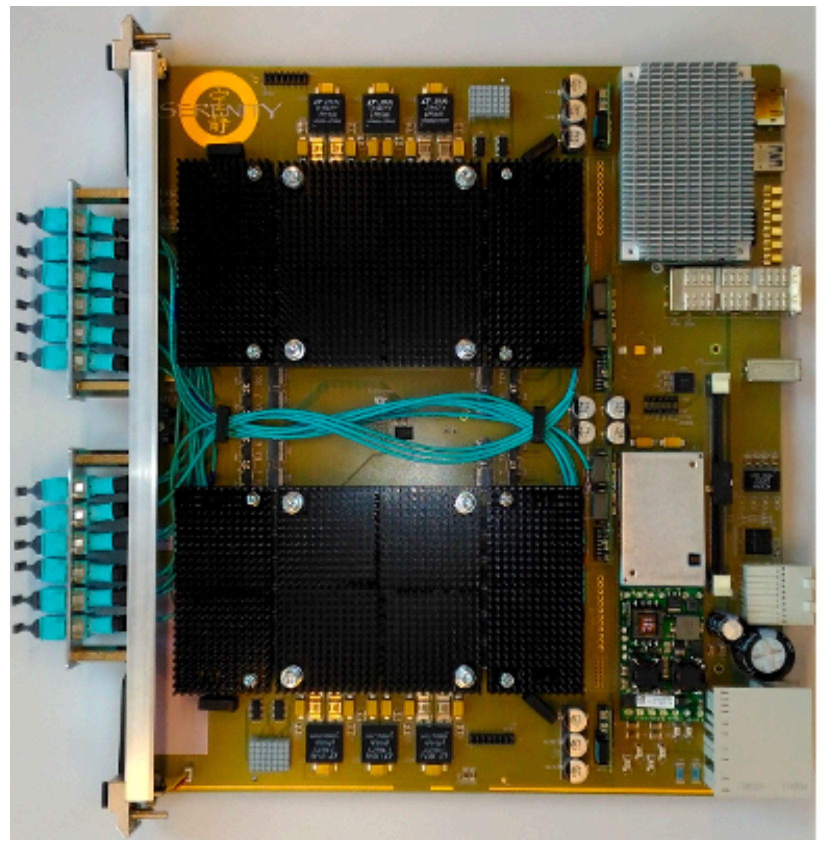

t = 0

Each small box is one Xilinx Ultrascale+ FPGA

t ≤ 12.5 μs

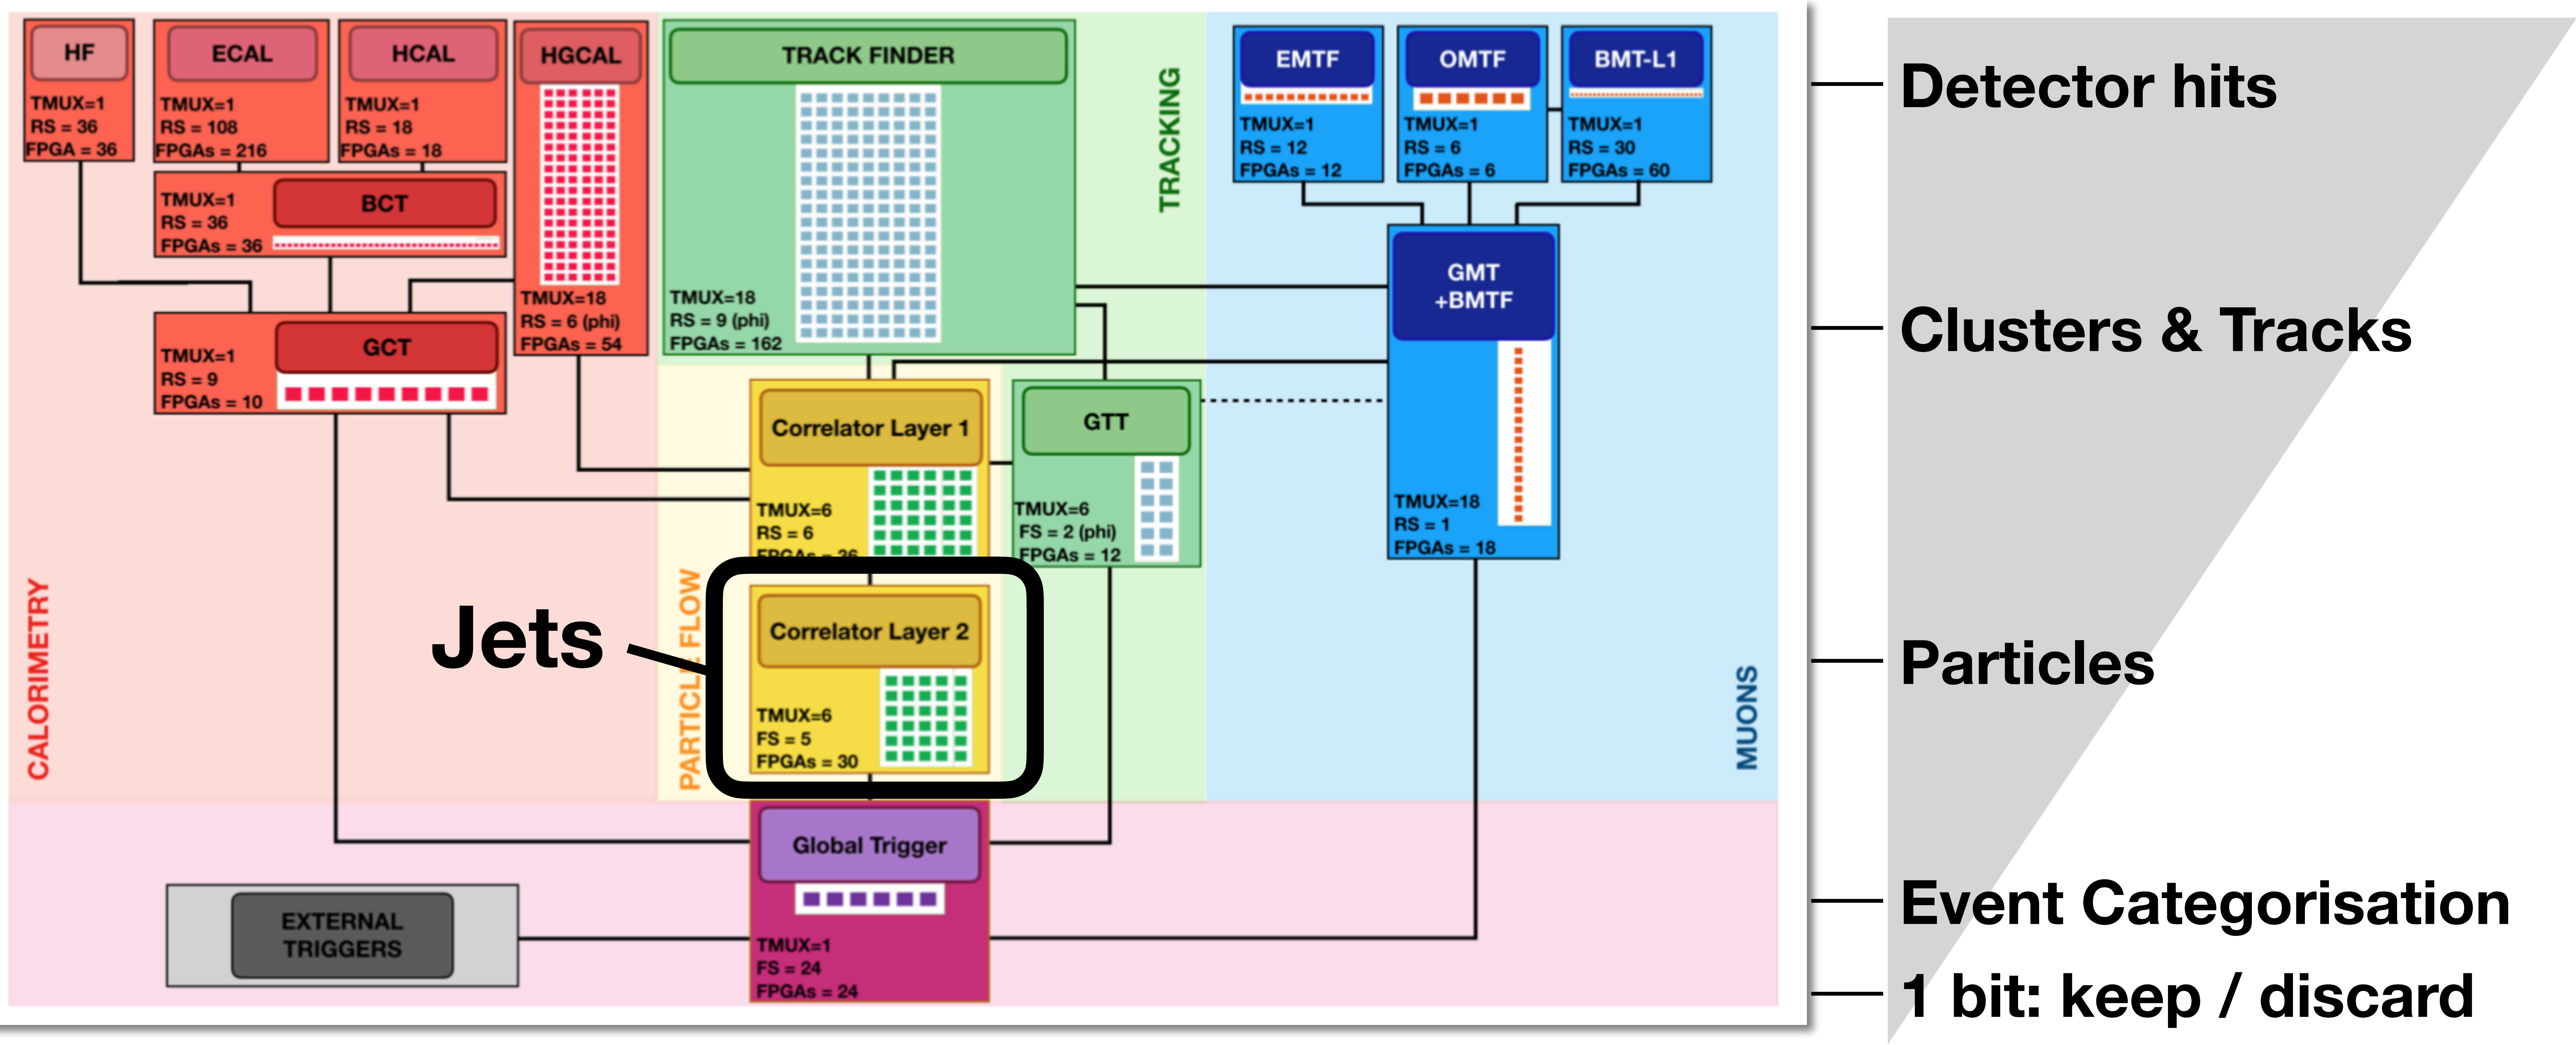

Deregionizer

- Data arrives into each Correlator Layer 2 FPGA over 30 optical fibres at 25 Gb/s from Correlator Layer 1 FPGAs
  - 64 bits per particle
- ‘Time Multiplex Factor 6’ architecture means data arrives over 6 LHC bunch crossings (150 ns)
  - 54 clock cycles at 360 MHz to receive one event
- Particles arrive grouped by regions in (η, ϕ) roughly 0.5 x 0.5
  - Packet structure shown by bold lines in graphic
  - Sorted by p_T within region
- Not very convenient to run a jet algorithm directly on this
- So ‘Deregionizer’ accumulates particles over the event into one flattened array in FPGA registers
  - Produces a new array of particles every 150 ns
- Truncate to 128 particles (justification in backup)

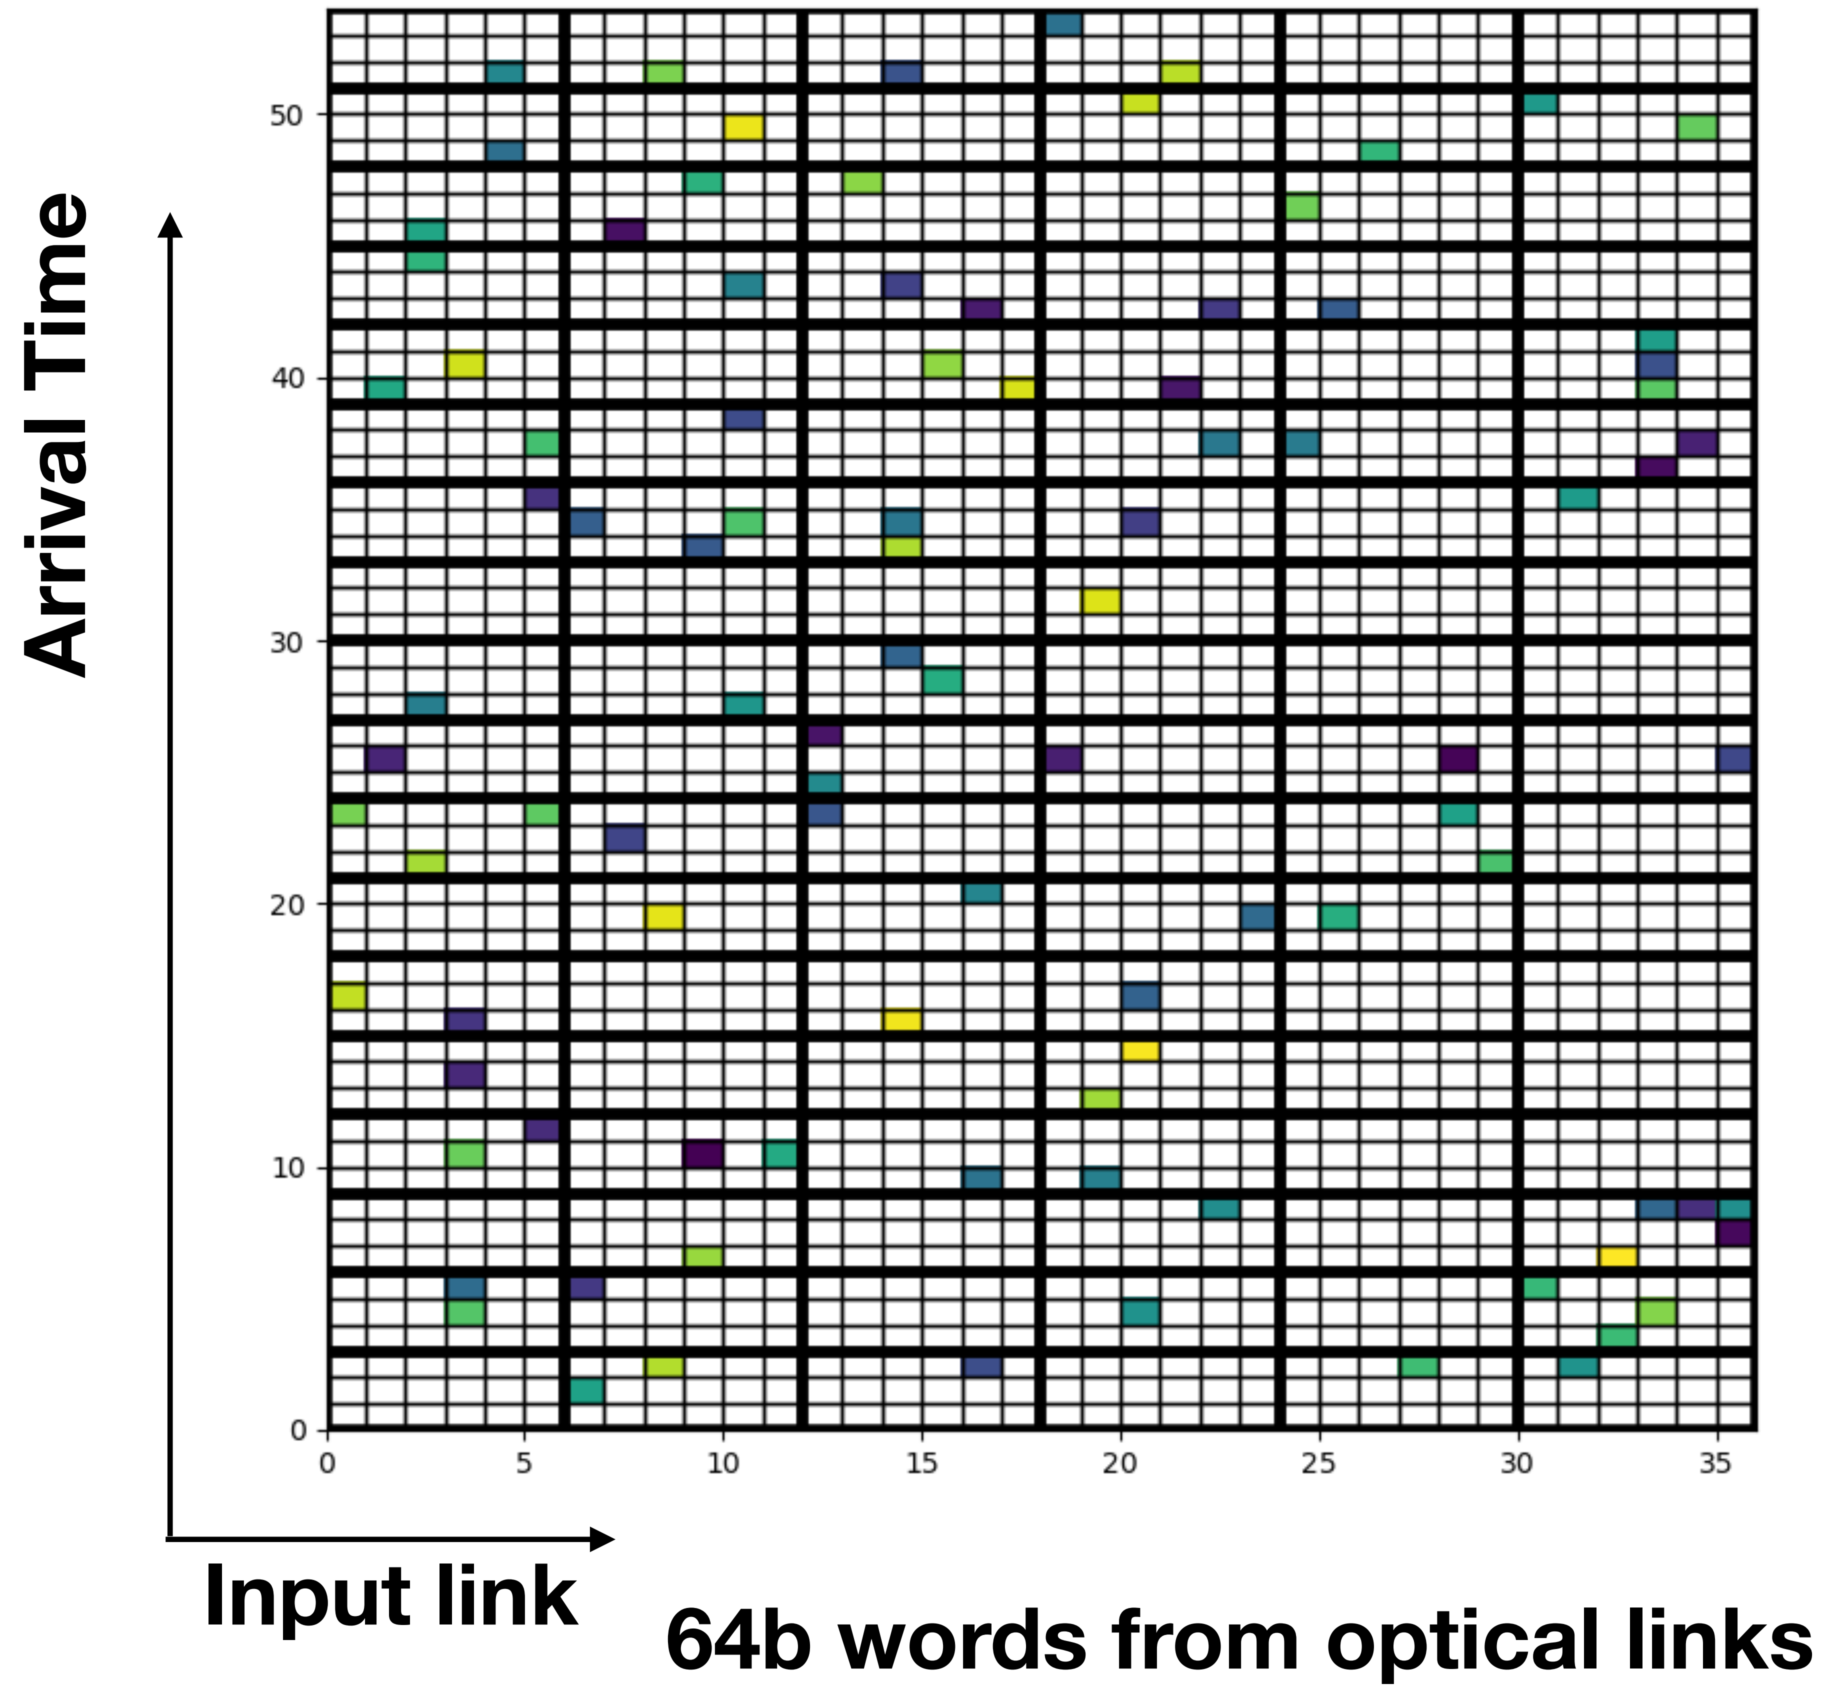

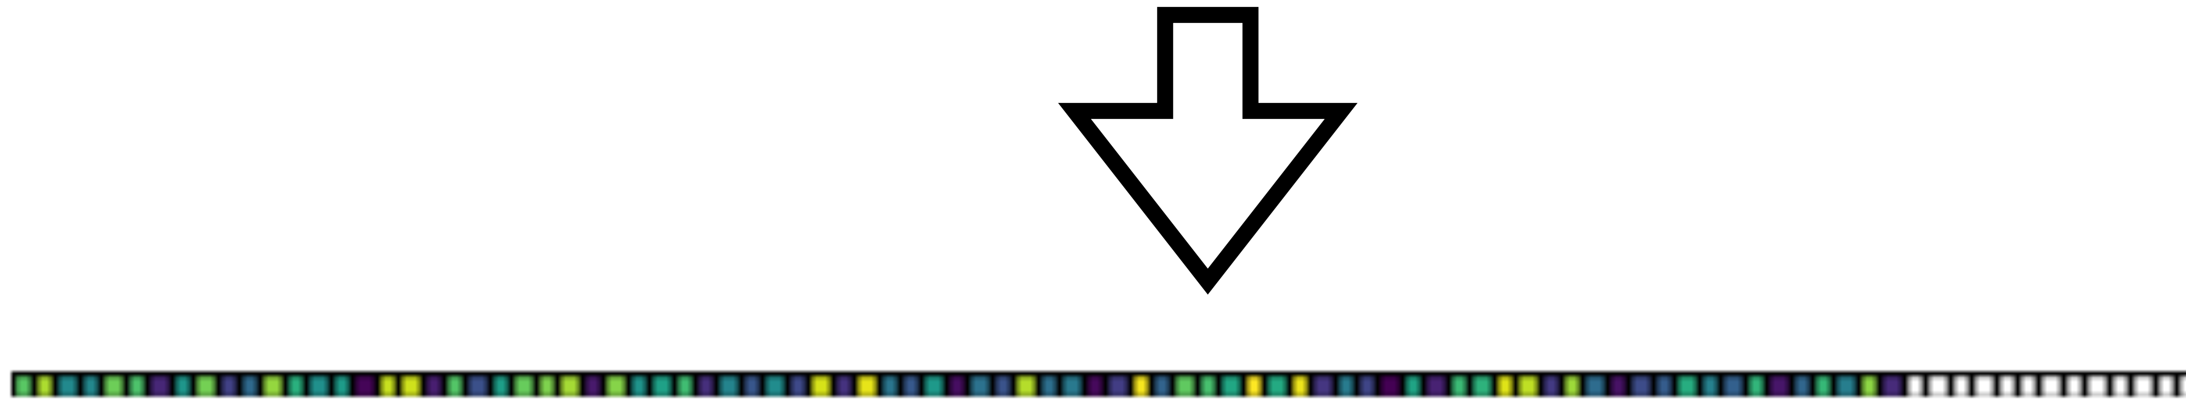

Seeded Cone Algorithm

1. Find the highest p_T particle as the seed for a jet
2. Compute distance between each particle and the seed
    - And whether in or out of jet radius
3. Compute jet axis from constituents, correct jet energy
4. Mask seed and constituents out from event, go back to 1

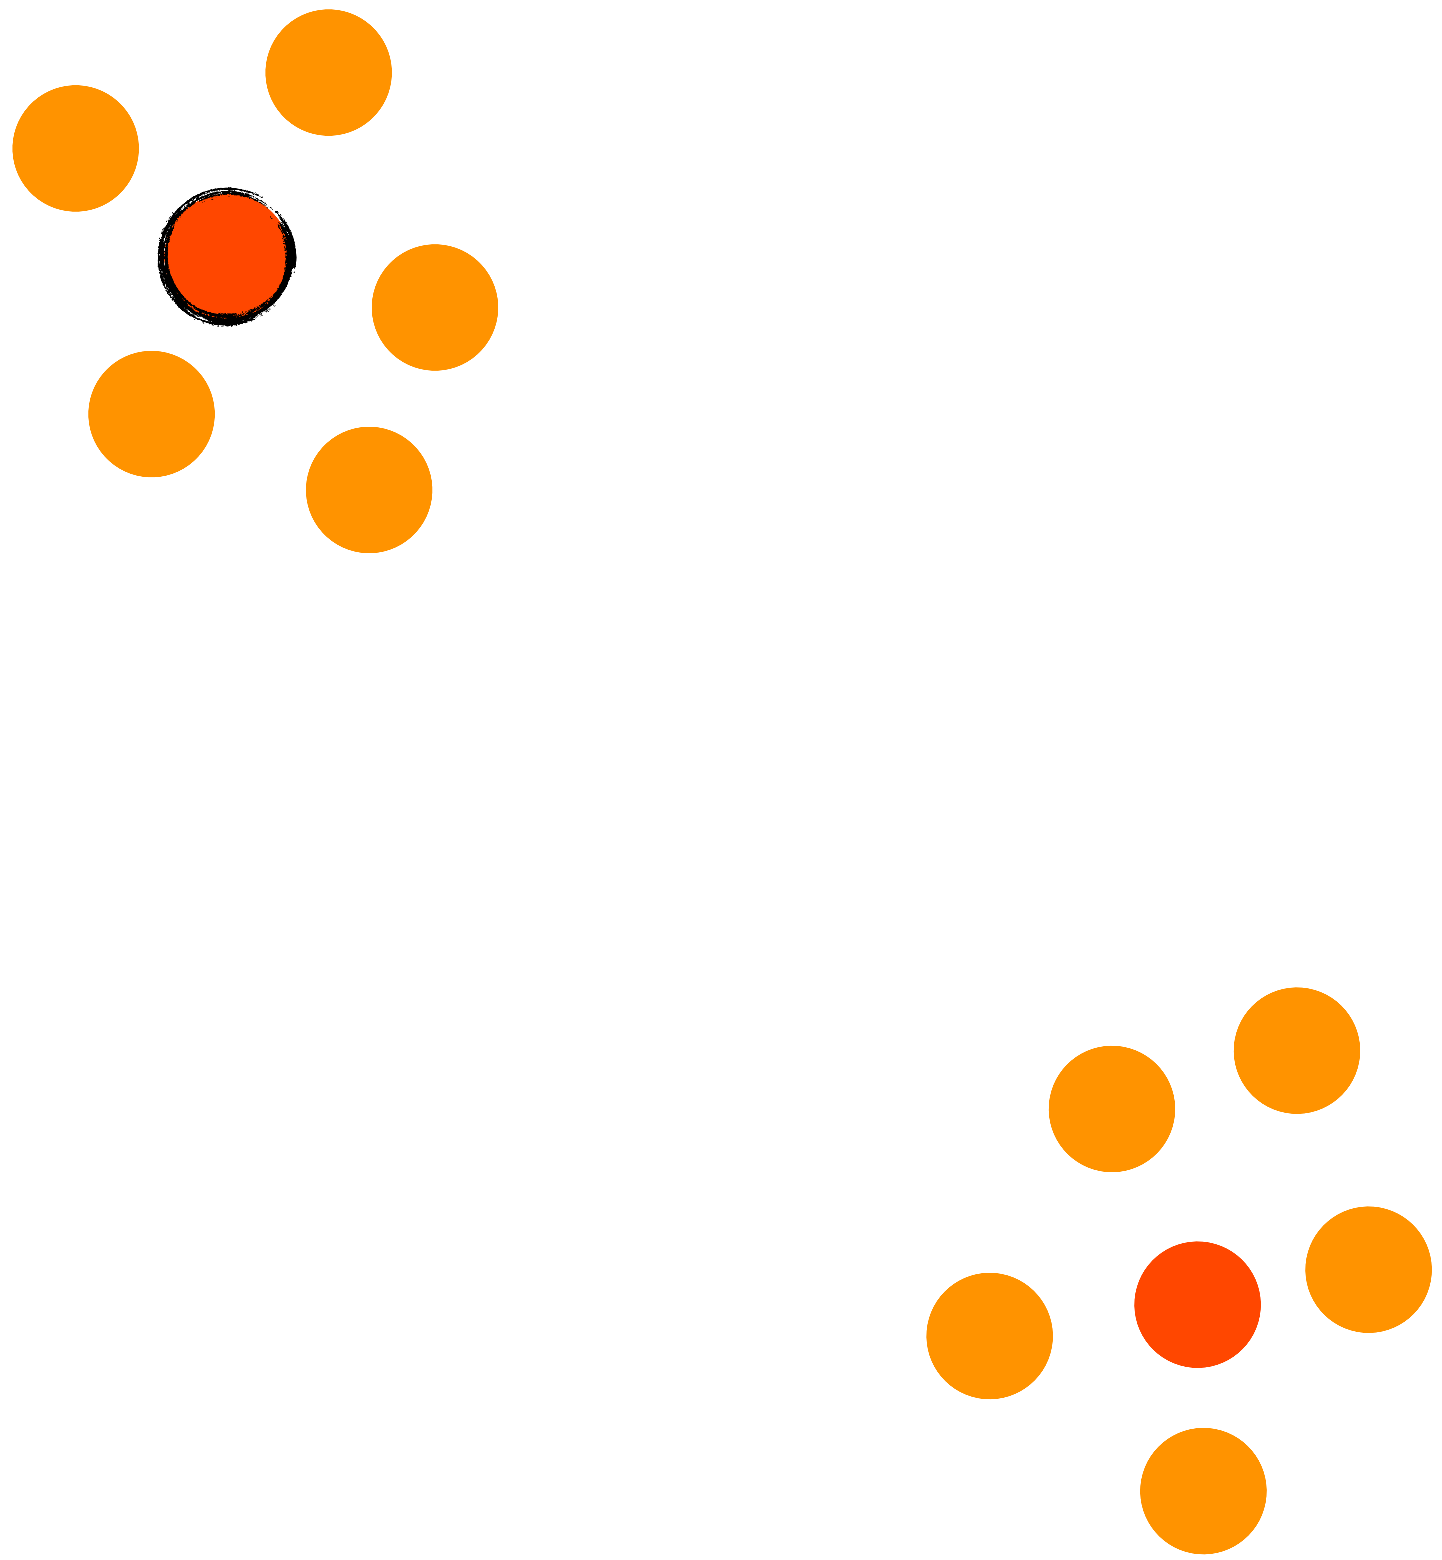

Seeded Cone Algorithm

1. Find the highest p_T particle as the seed for a jet
2. Compute distance between each particle and the seed
    - And whether in or out of jet radius
3. Compute jet axis from constituents, correct jet energy
4. Mask seed and constituents out from event, go back to 1

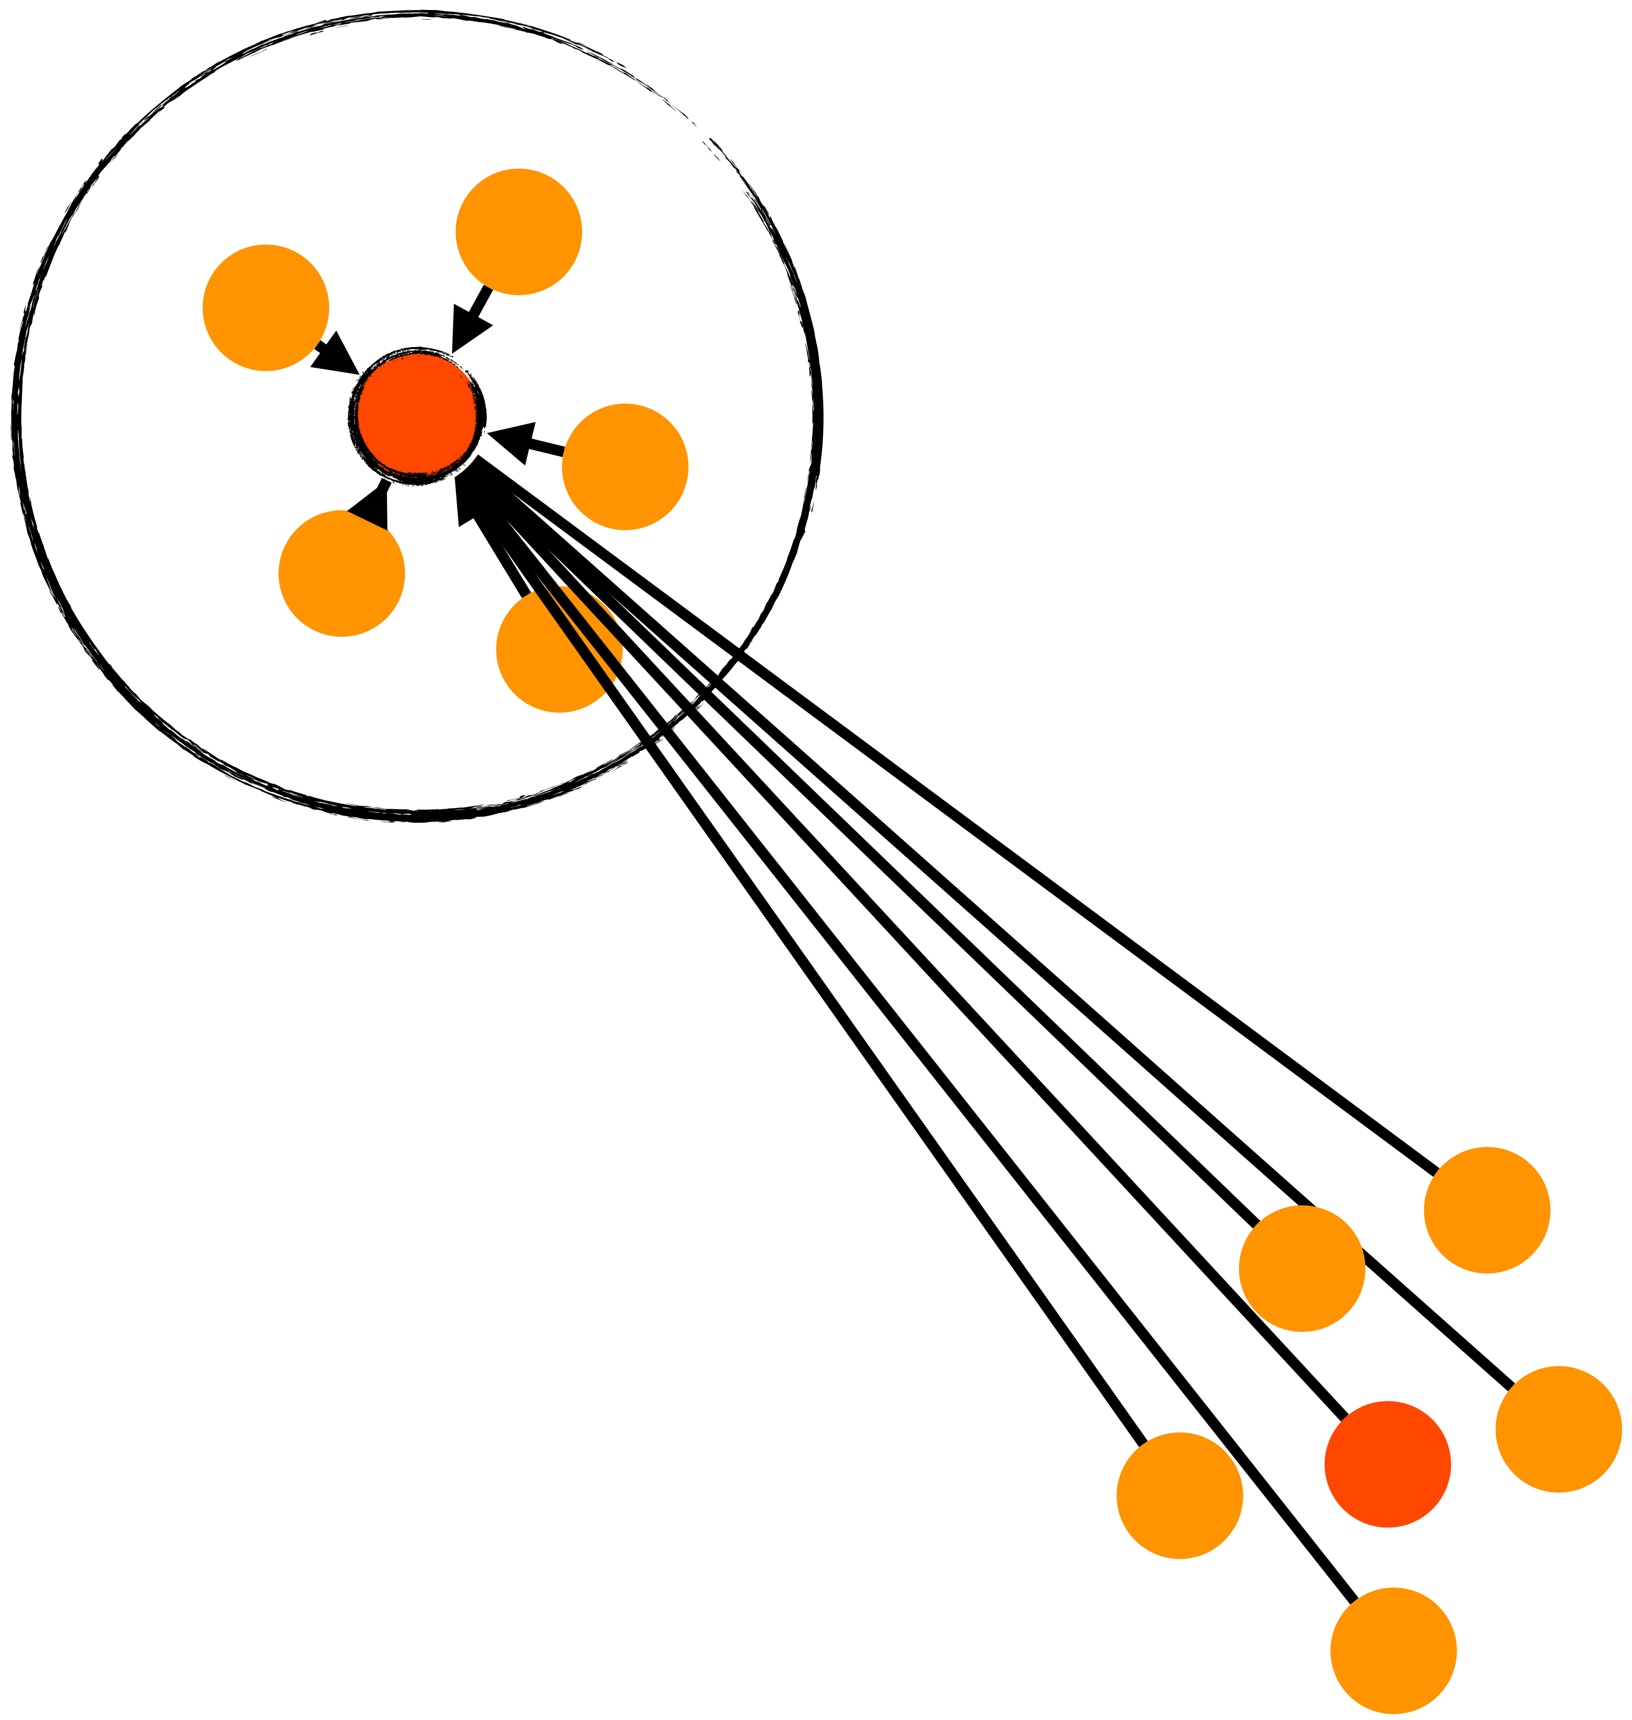

Seeded Cone Algorithm

1. Find the highest p_T particle as the seed for a jet
2. Compute distance between each particle and the seed
   - And whether in or out of jet radius
3. Compute jet axis from constituents, correct jet energy
4. Mask seed and constituents out from event, go back to 1

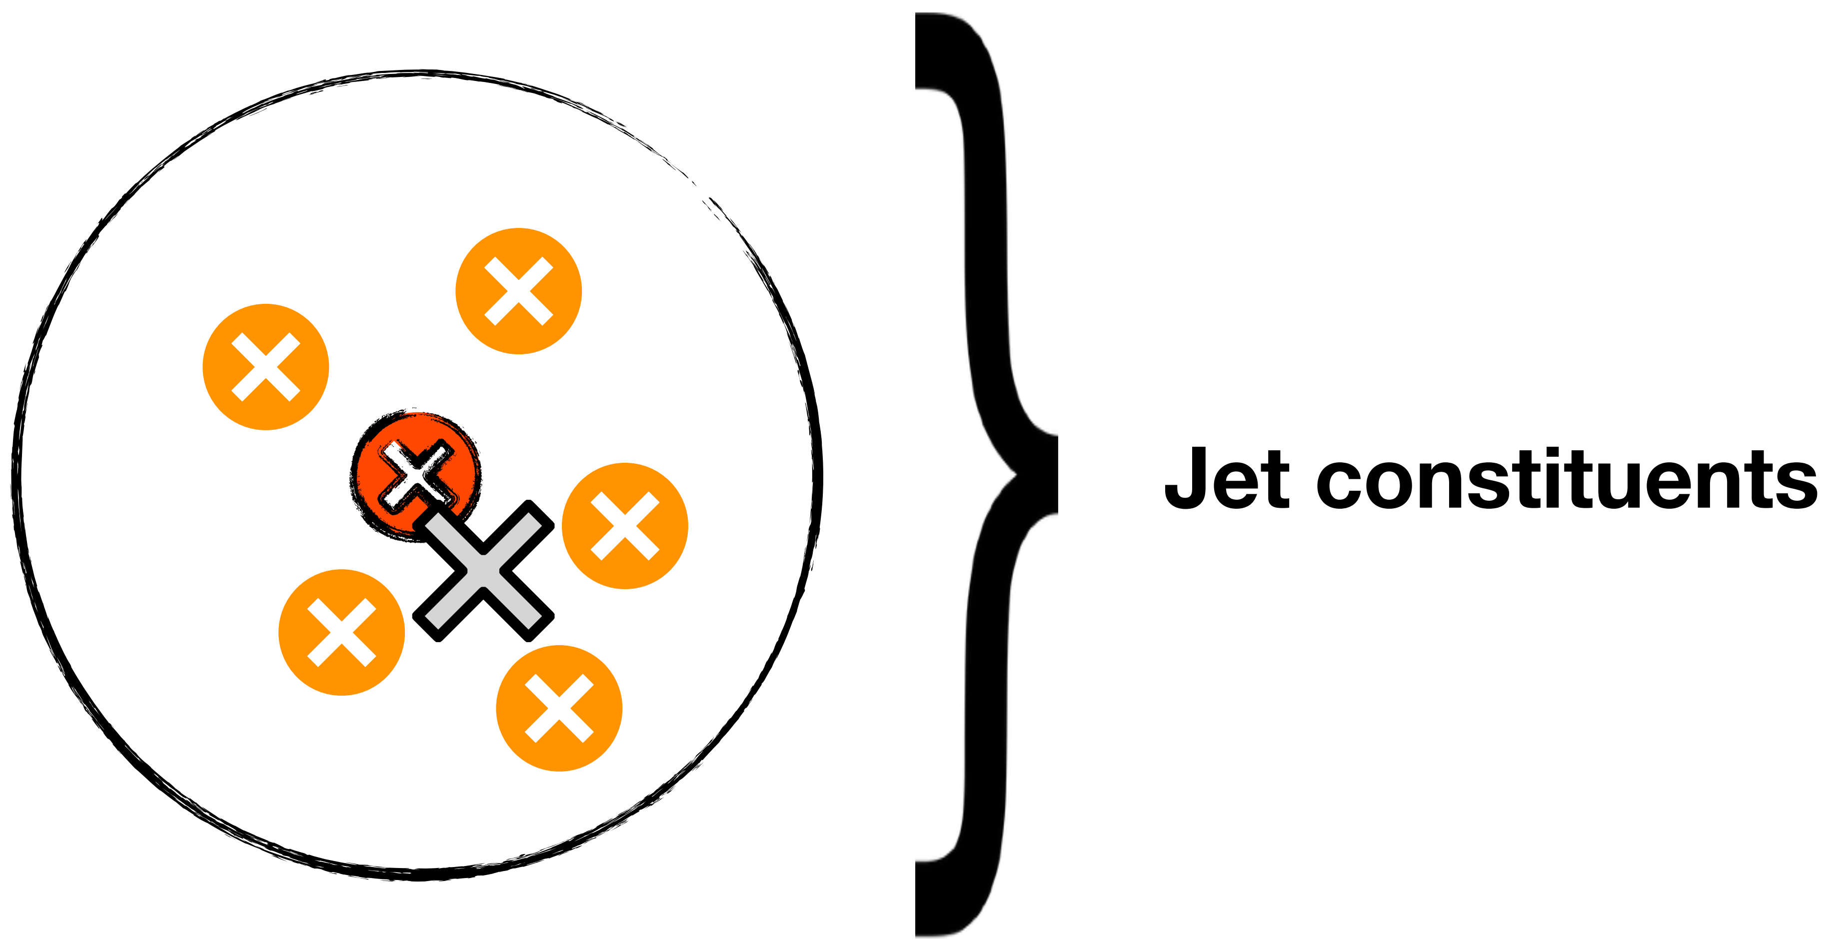

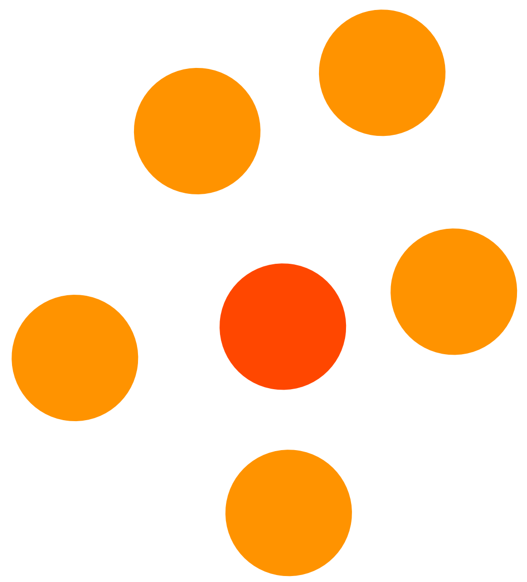

Seeded Cone Algorithm

1. Find the highest p_T particle as the seed for a jet
2. Compute distance between each particle and the seed
    - And whether in or out of jet radius
3. Compute jet axis from constituents, correct jet energy
4. Mask seed and constituents out from event, go back to 1

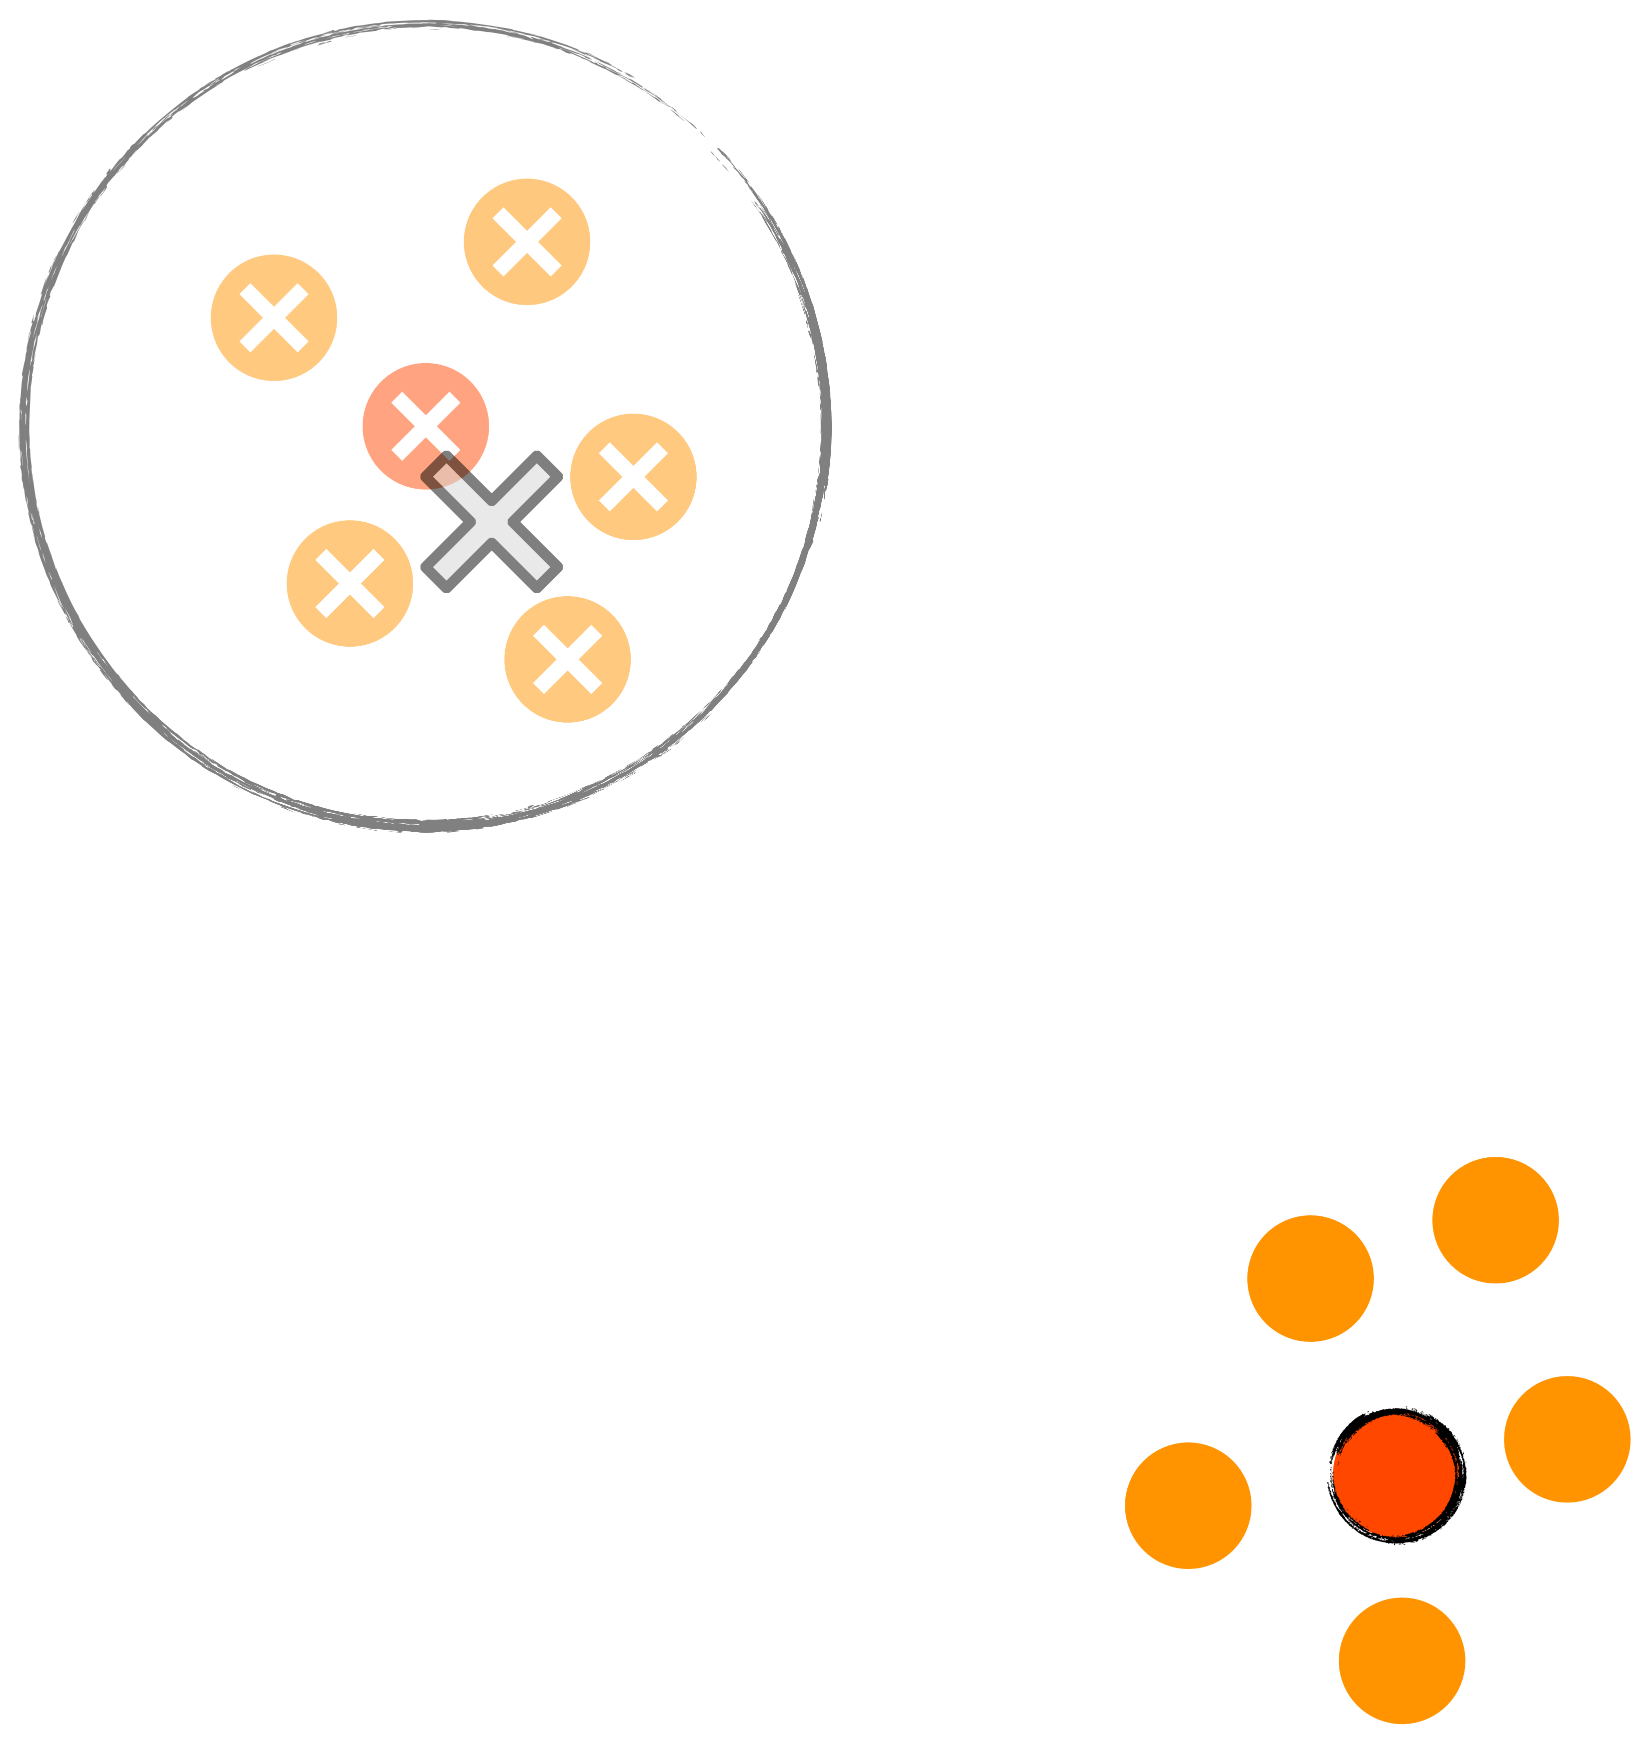

Seeded Cone Algorithm

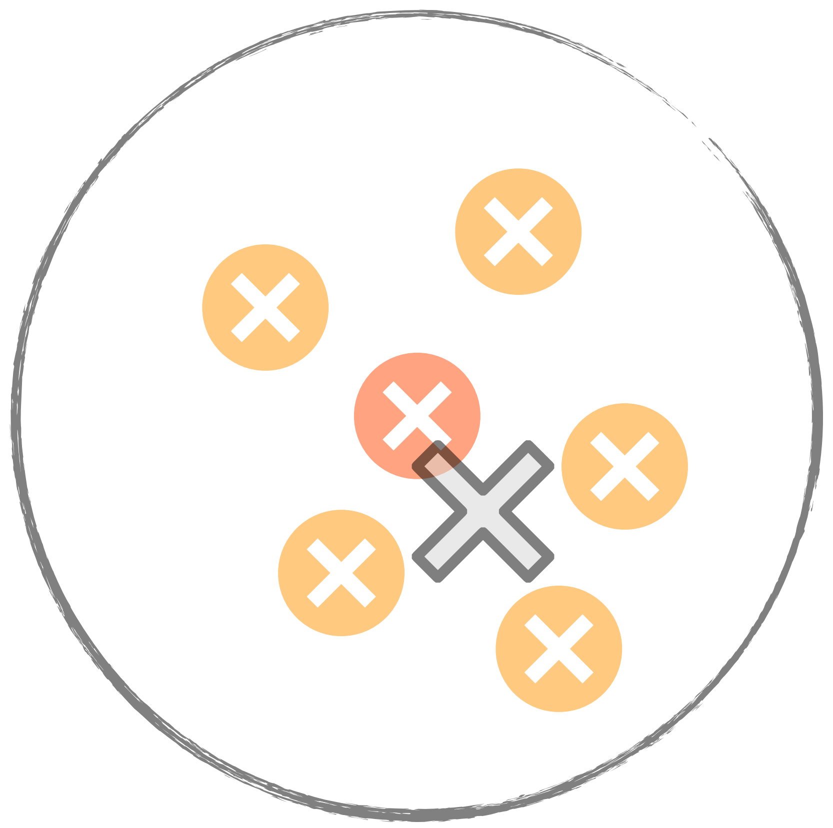

1. Find the highest p_T particle as the seed for a jet
2. Compute distance between each particle and the seed
   - And whether in or out of jet radius
3. Compute jet axis from constituents, correct jet energy
4. Mask seed and constituents out from event, go back to 1

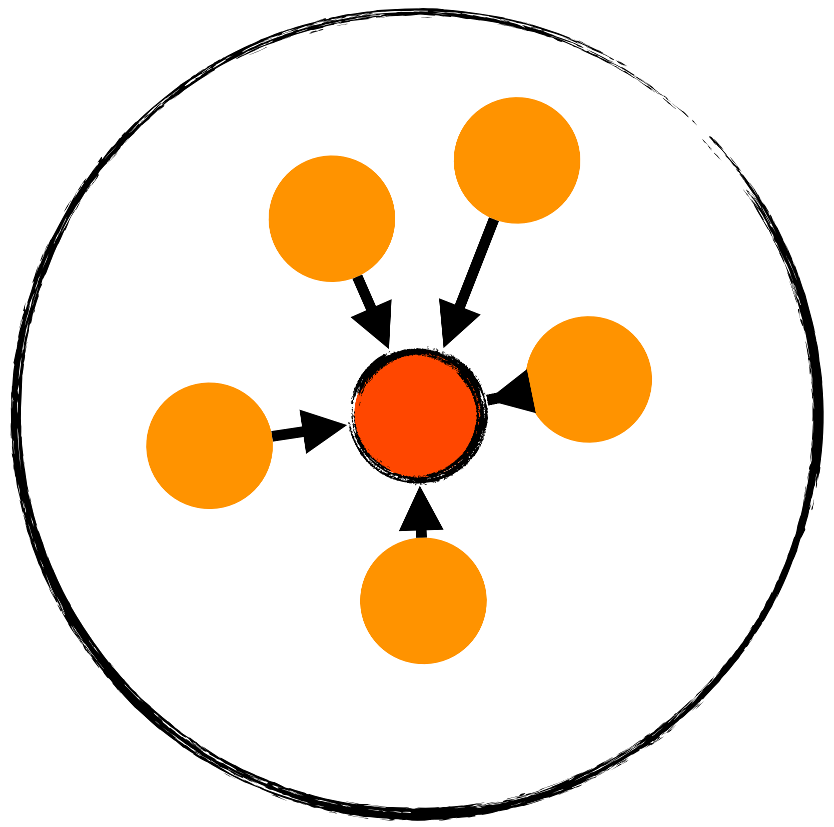

Seeded Cone Algorithm

1. Find the highest p_T particle as the seed for a jet
2. Compute distance between each particle and the seed
    - And whether in or out of jet radius
3. Compute jet axis from constituents, correct jet energy
4. Mask seed and constituents out from event, go back to 1

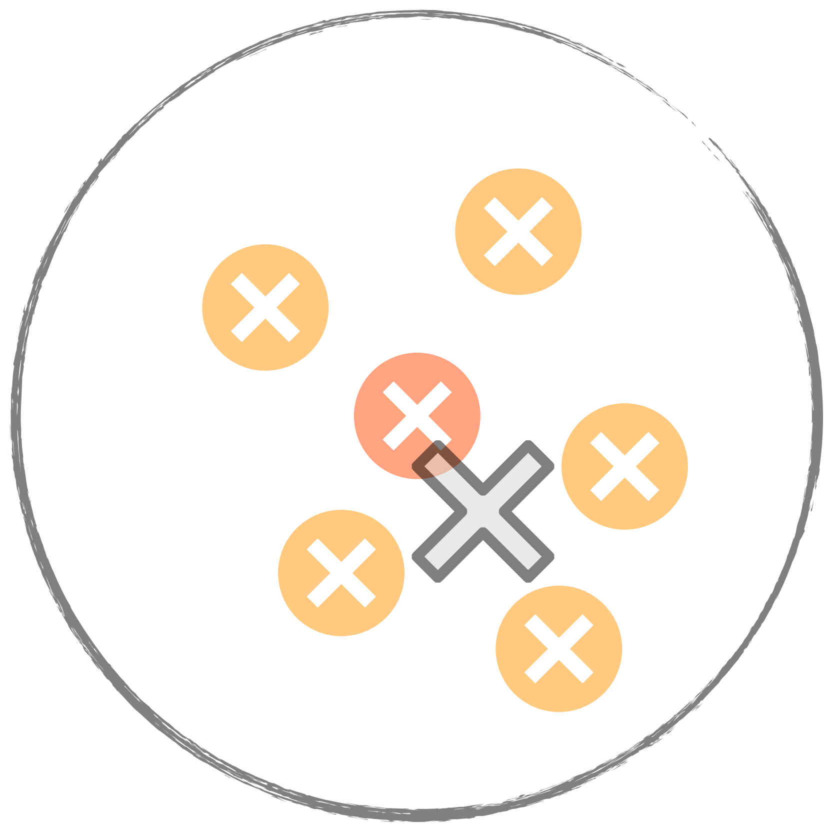

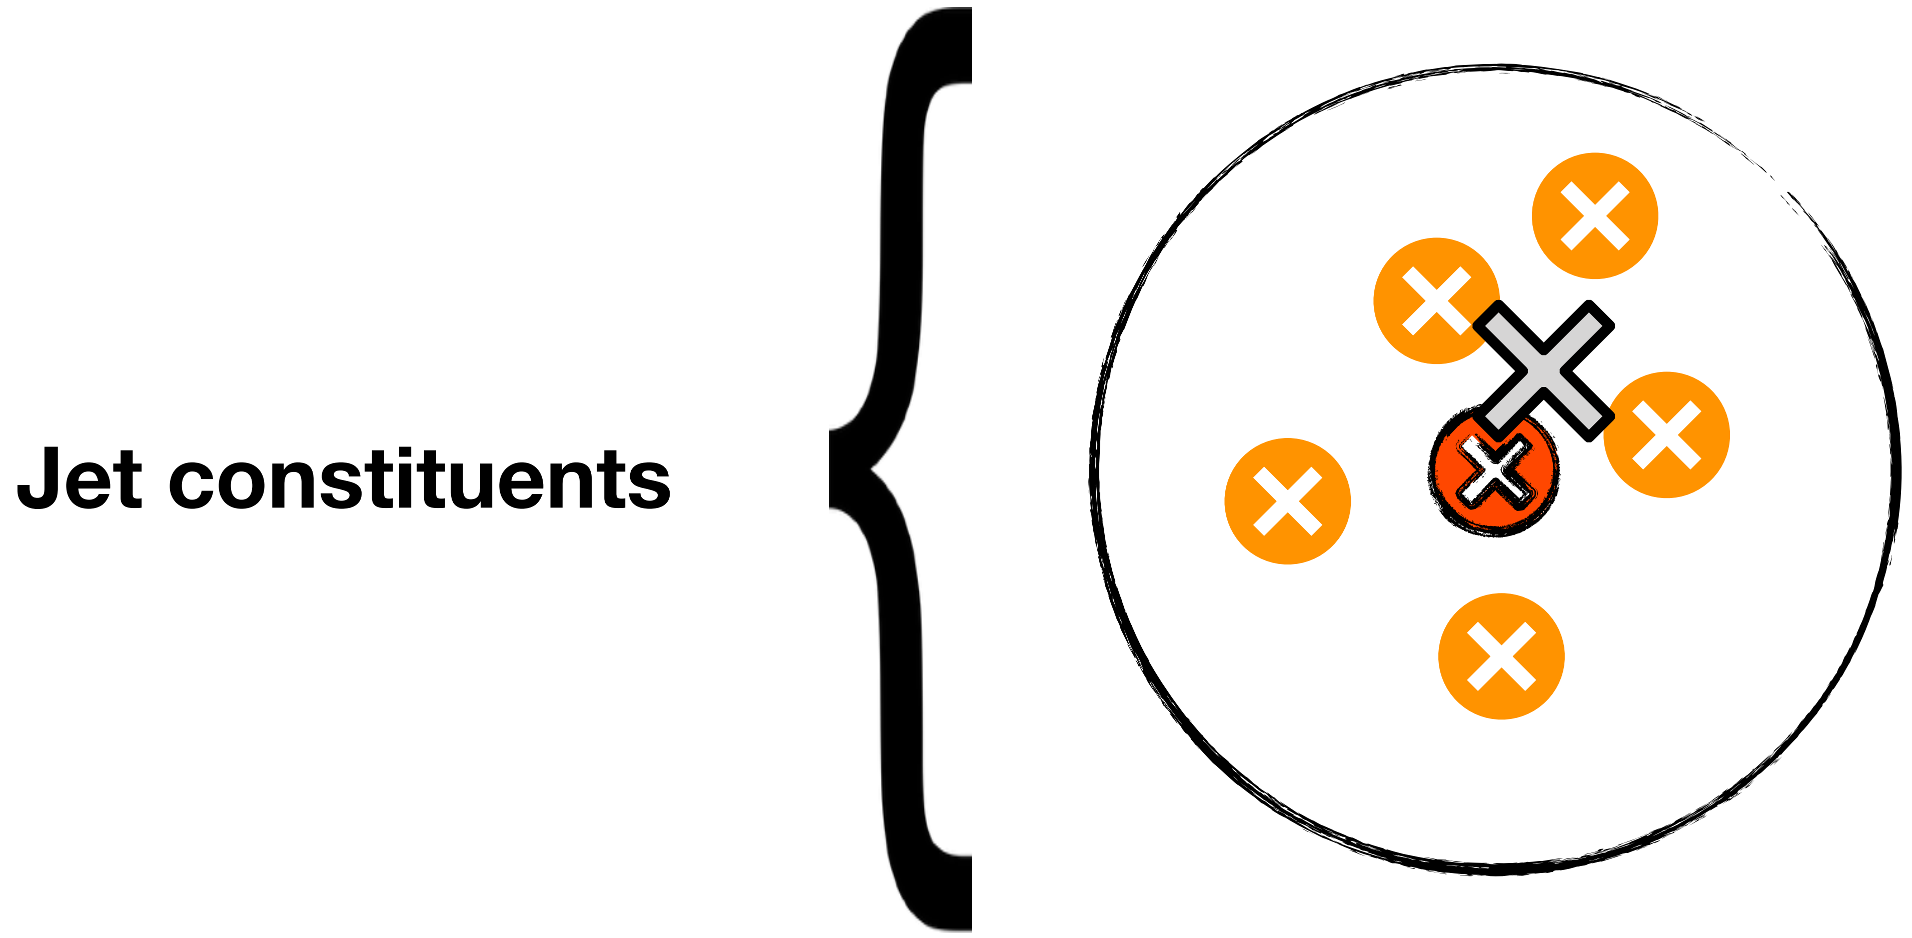

Implementation

colours represent different events

- Target Serenity ATCA card with Xilinx VU9P FPGA
- Input particles are in registers so can read all simultaneously
  - Parallelise as much as possible e.g. all ΔR computations
- Uses High Level Synthesis (HLS) for arithmetic parts
  - FPGA-optimized C++
- Algorithm has a “loop carried dependency” - iteration `i` depends on iteration `i-1`
  - Need to find constituents of jet `i-1` before finding seed of jet `i`
- Uses VHDL for optimal looping / control
  - Utilise partially-filled loop pipeline to simultaneously process different jet parameters (R=0.4, R=0.8) and events (up to 3)
- Split algorithm into modules for each task
  - Make the constituent finding loop task as short as possible

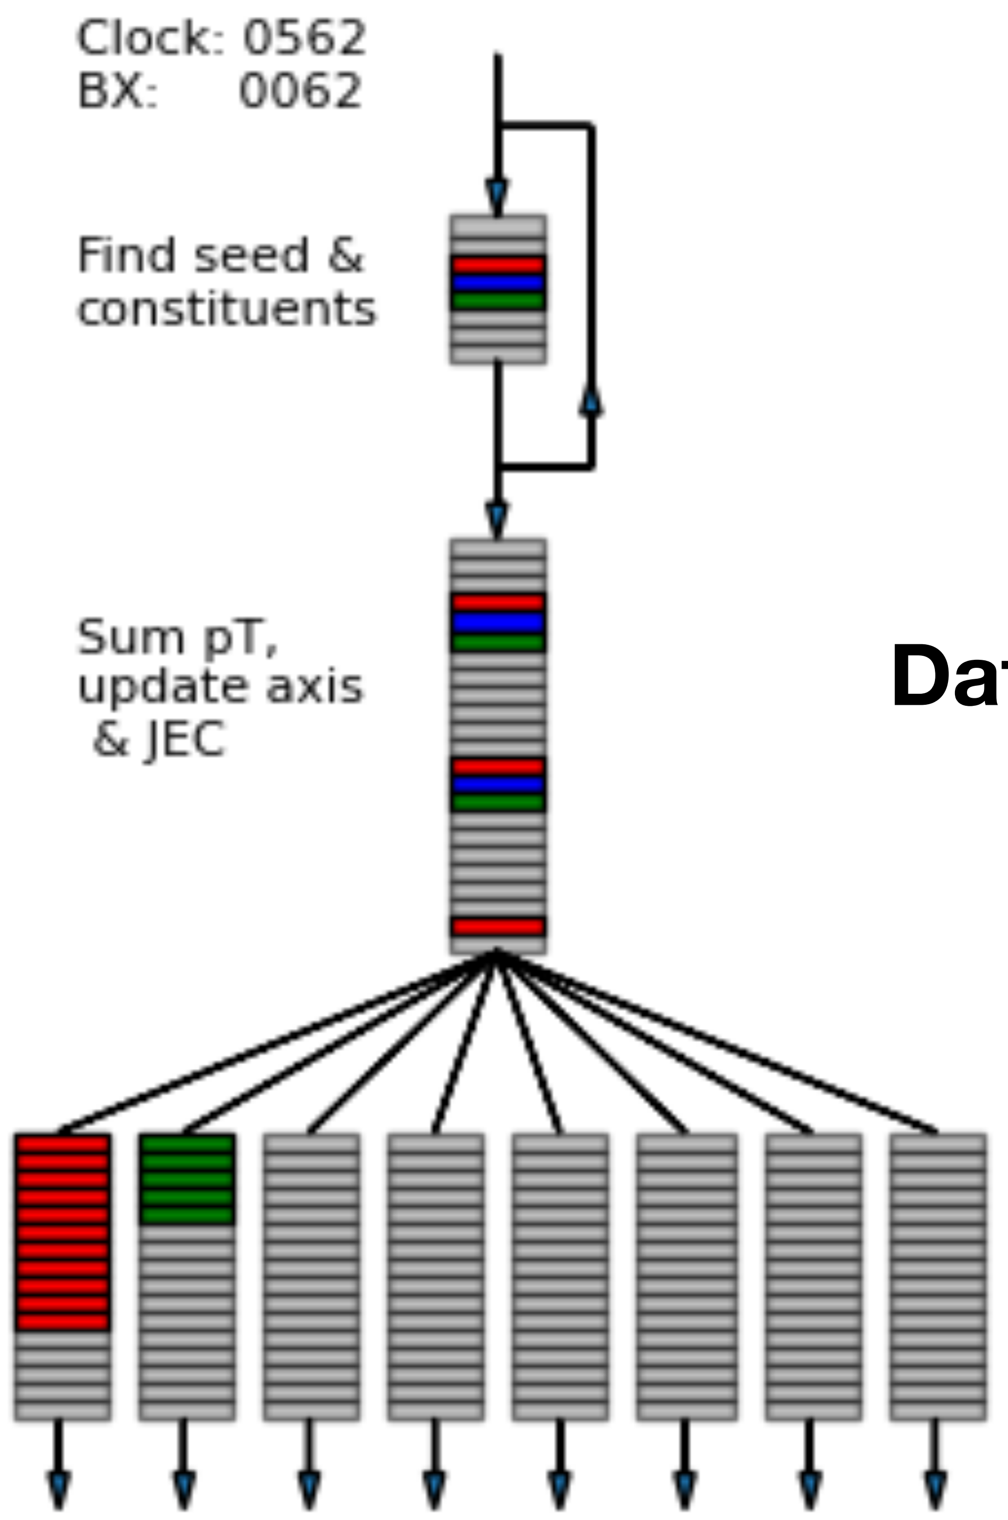

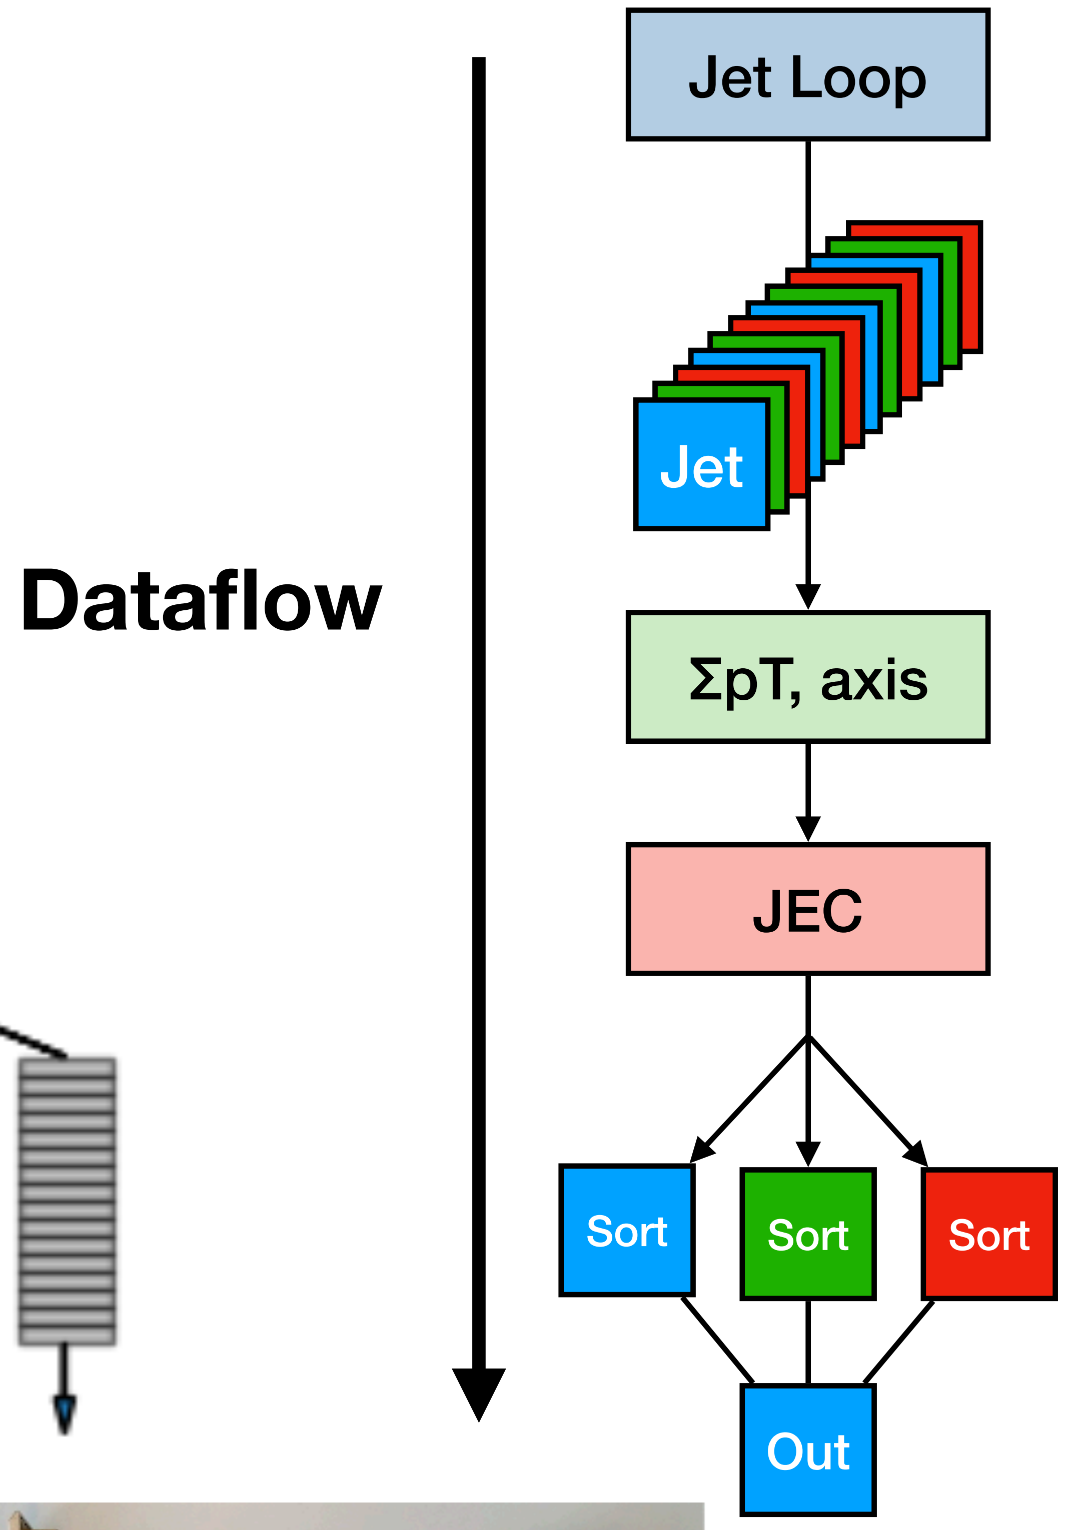

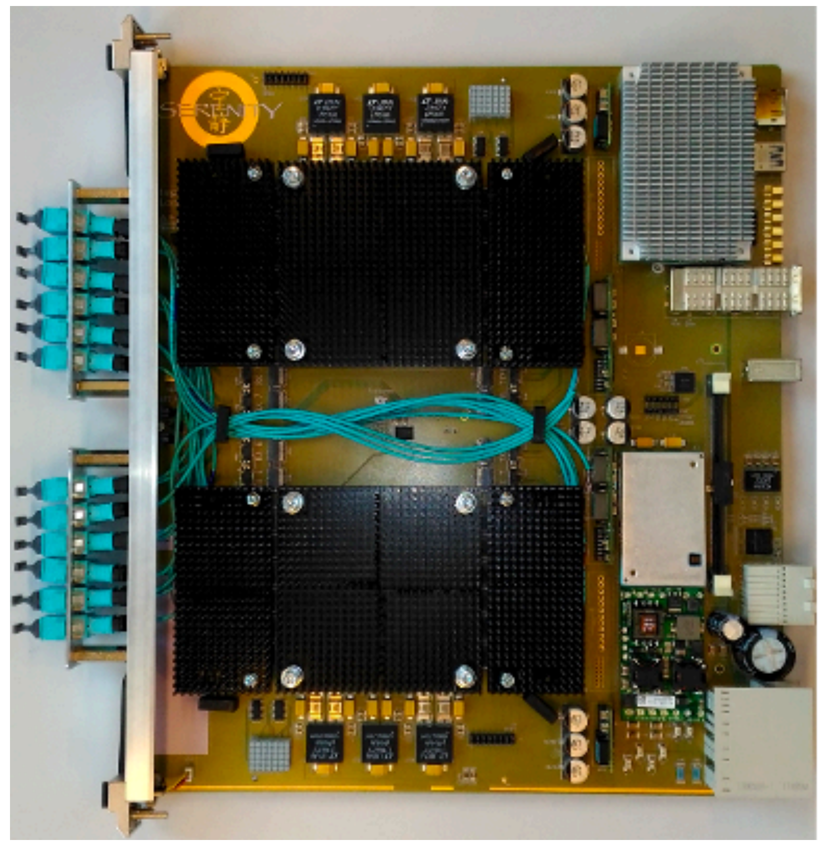

Latency

- Latency is dominated by latency of Jet Loop module (find next seed and jet constituents)
  - Optimized down to 8 clock cycles per iteration (22.2 ns at 360 MHz)
  - Find 16 jets for robustness in high multiplicity events - some seeds might turn out to be isolated energetic particles
- First particle in to last jet out: 744 ns
- Events arrive every 150 ns and are processed in the same logic modules thanks to pipelined implementation
- Second jet collection (R=0.8) produced exactly one clock cycle after the first (2.78 ns)
- 16 jets per event / 1 event per 150 ns → more than 100 million jets per second per FPGA

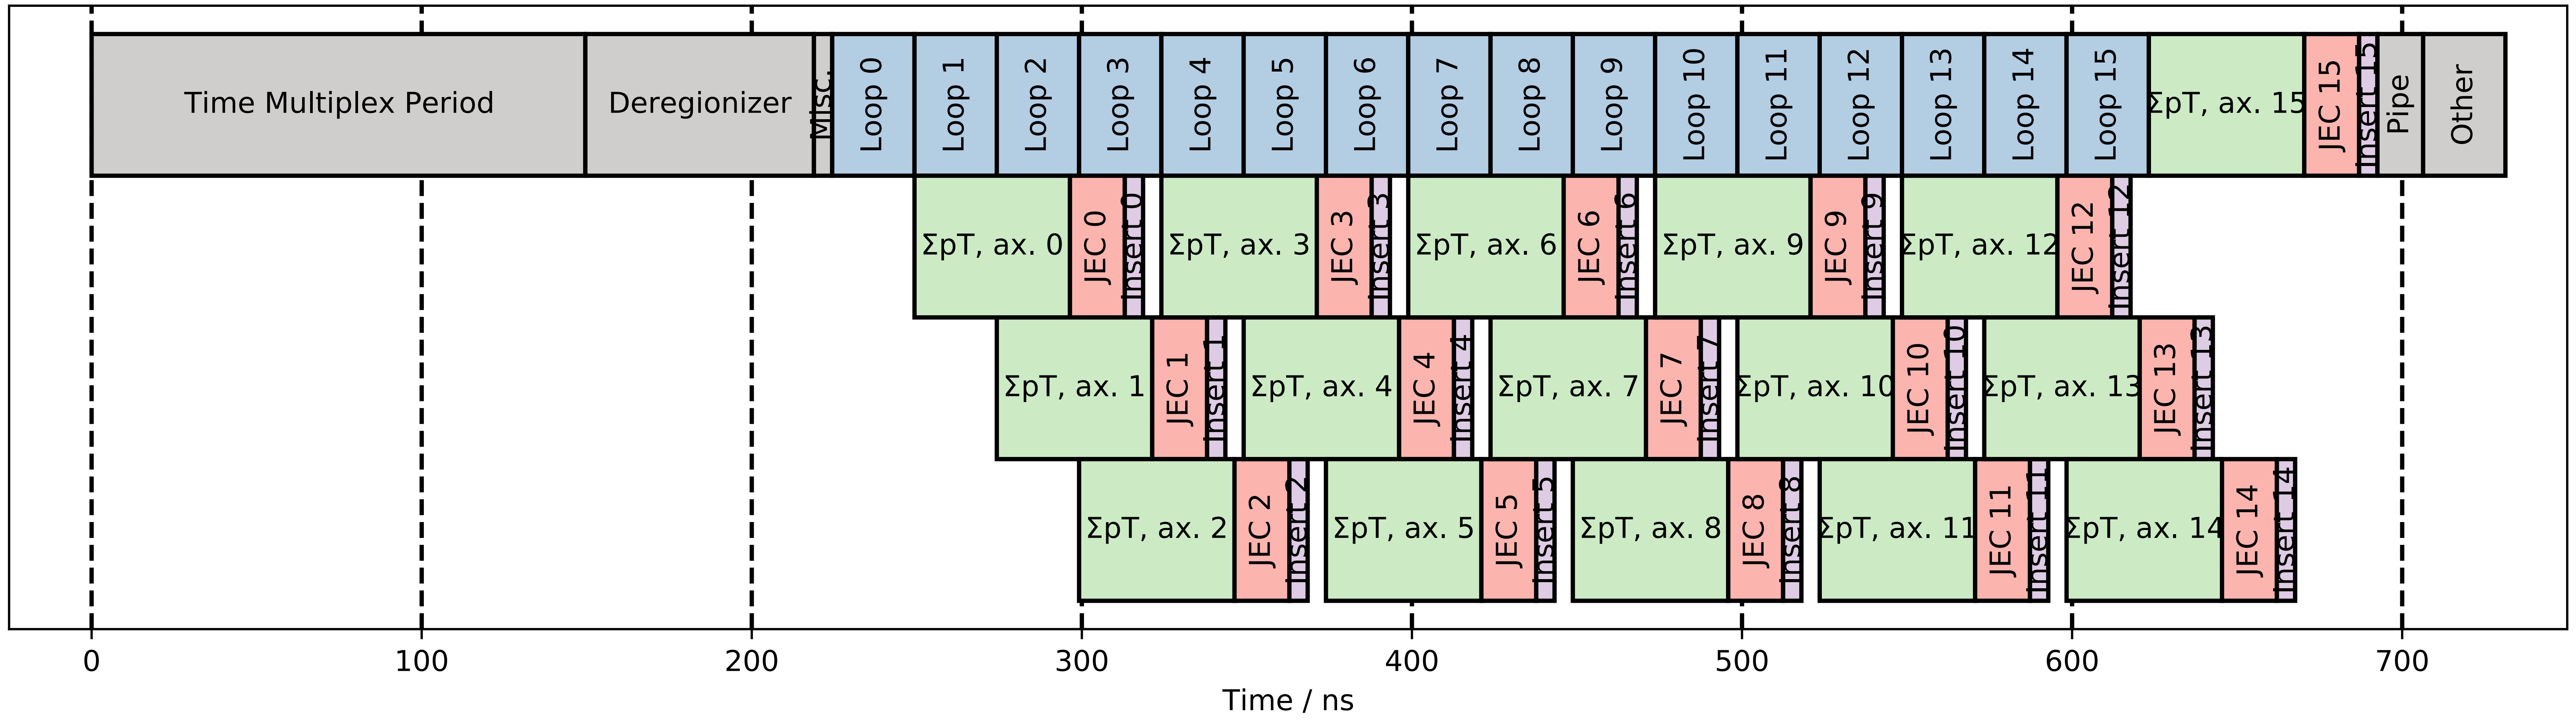

Implementation Results

- Design fits FPGA and meets timing constraints at 360 MHz in Xilinx VU9P FPGA
- Initial placement constraints from I/O locations, data flows top-to-bottom
- Design is quite congested (fan-in, fan-out) so spread out modules as much as possible
- Bottom Super Logic Region (SLR) so far unused, anticipating addition of jet tagging NNs
- Hardware performance validated against a C++ emulator in CMS software framework

Module	LUT %	FF %	DSP %	BRAM %
Total	7.41	6.17	7.62	0.07
Deregionizer	3.47	2.31	-	-
Jet Loop + Σp_T & axis + Sort	3.50	3.62	7.50	0.02 [0.5]
JECs	~0 [81]	~0 [109]	0.03 [2]	0.02 [0.5]
HT & MHT	0.28	0.12	0.07 [5]	0.02 [0.5]

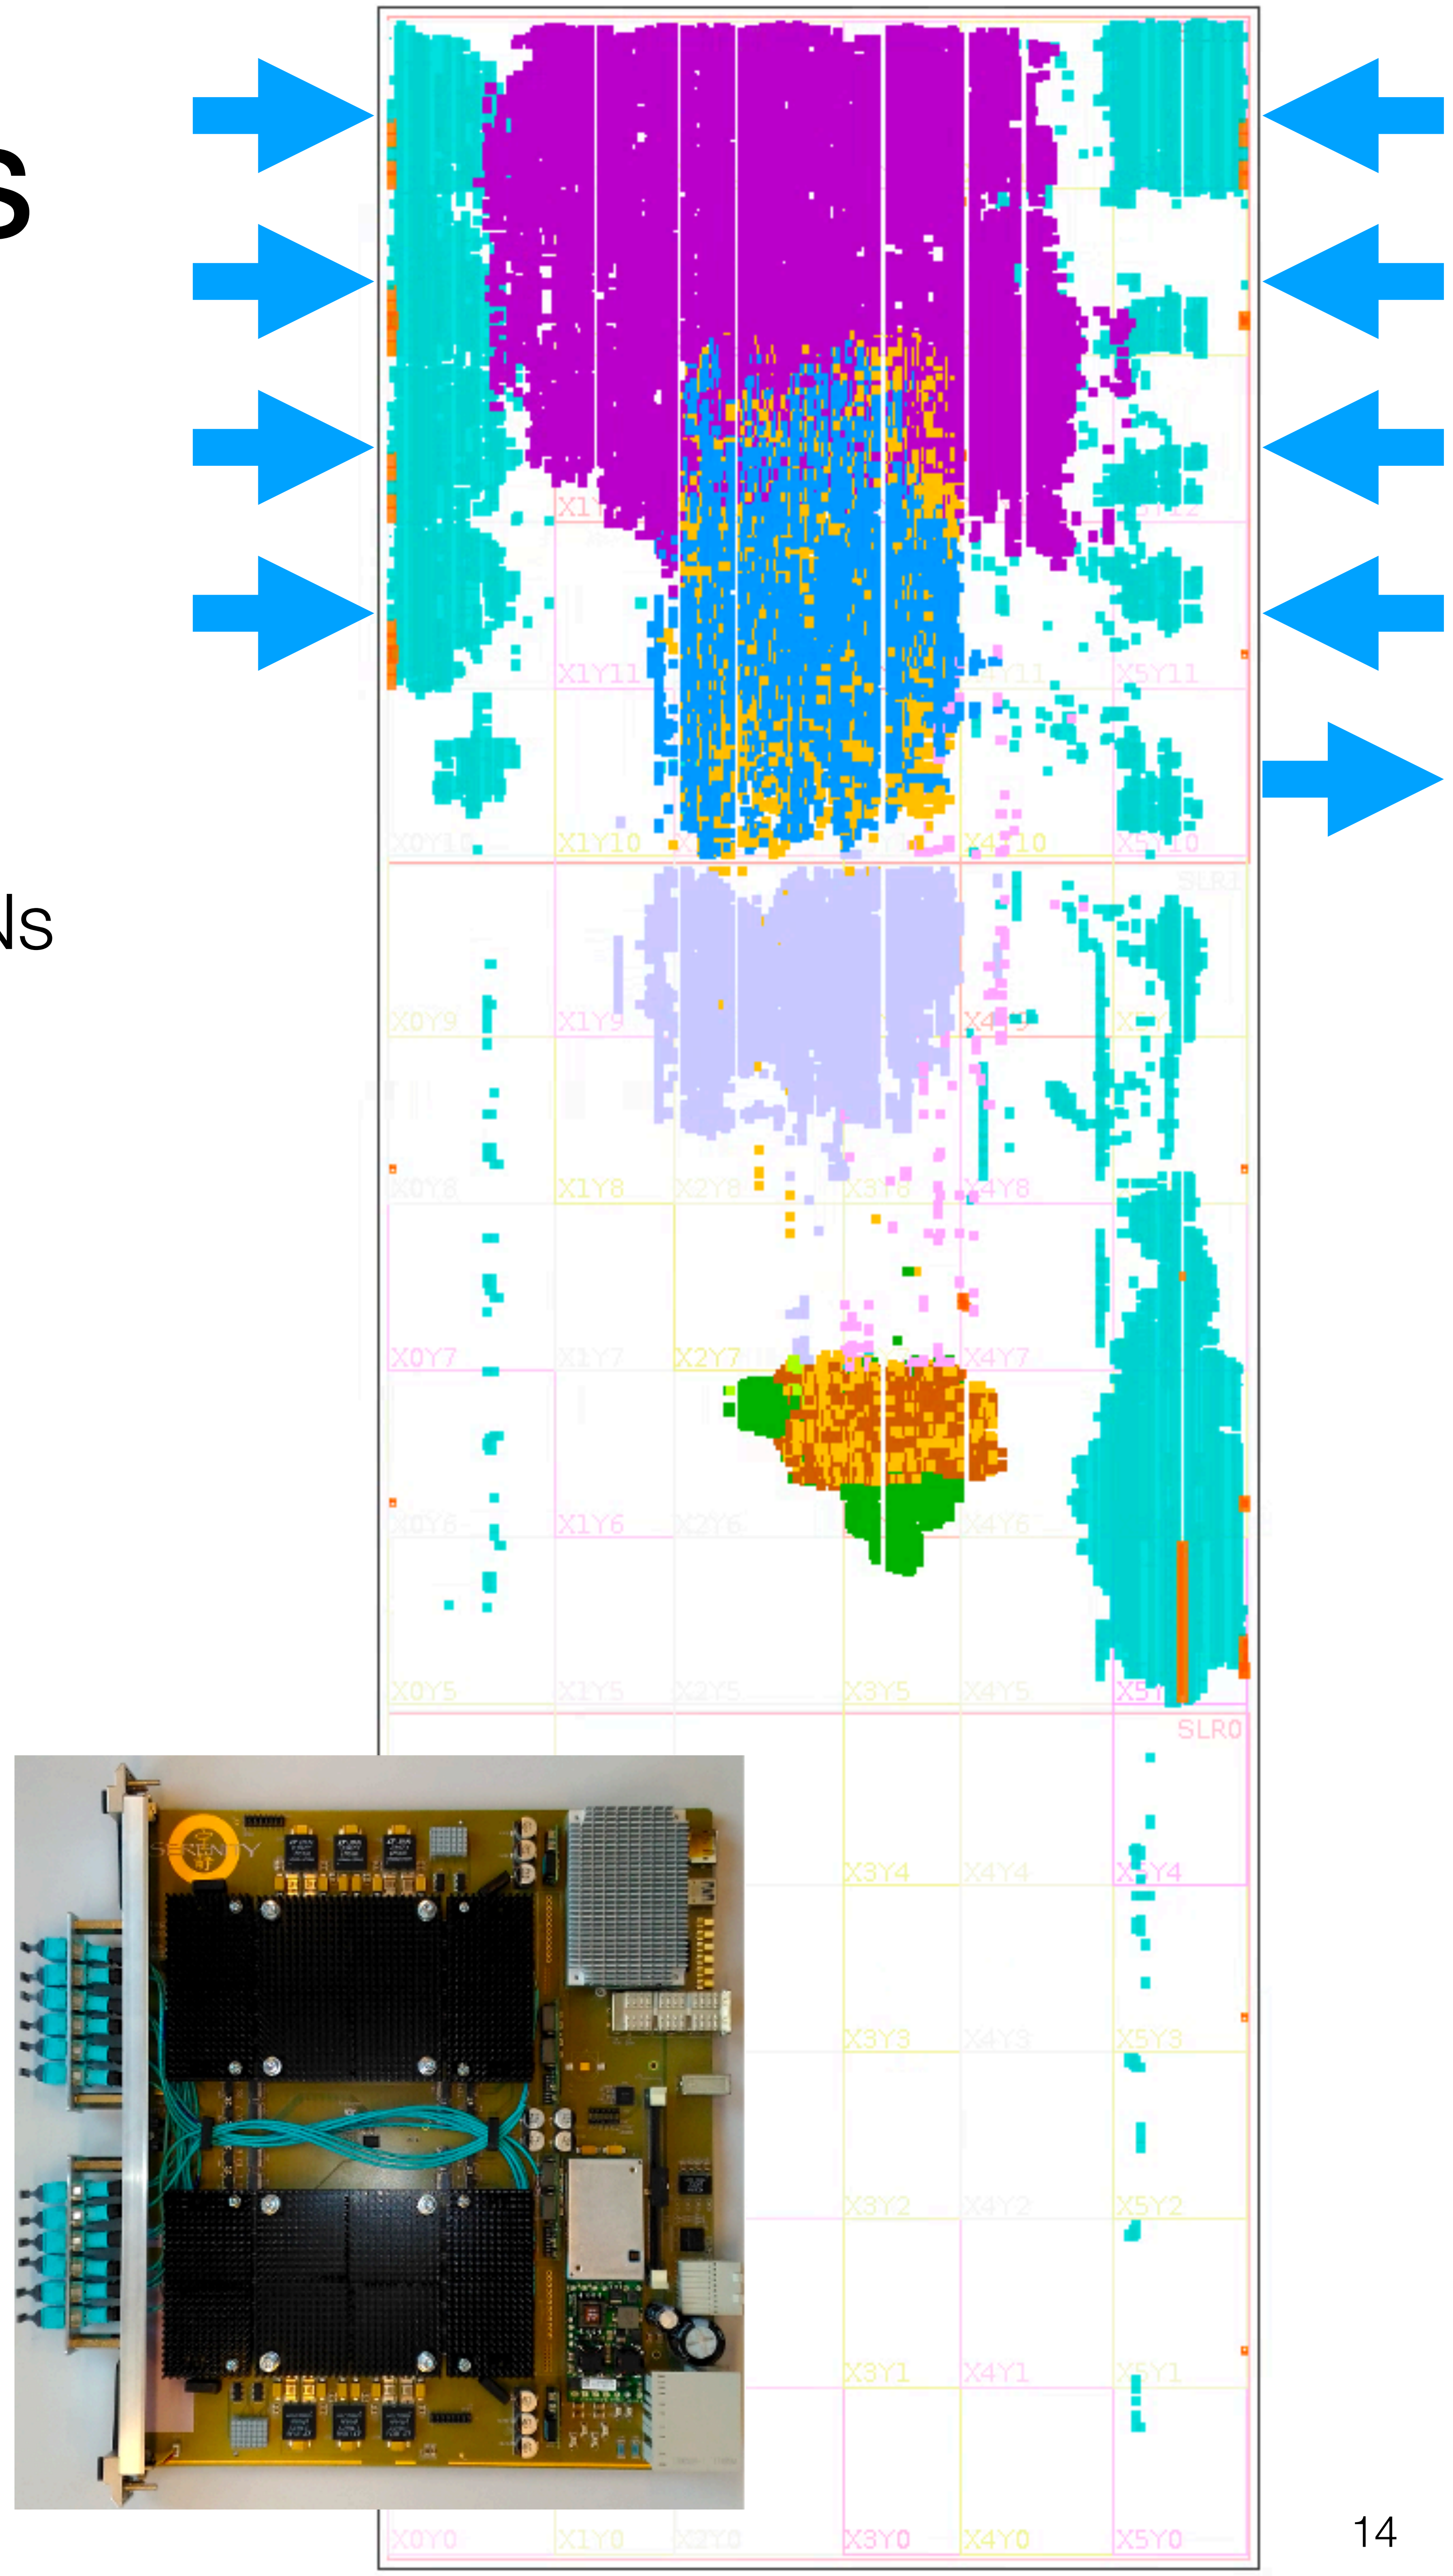

Performance

- Simulated events with tt + 200 pileup (top), QCD + 200 pileup (bottom) run through CMS detector simulation and L1T algorithm emulation
- Run Seeded Cone anti-k_t jet reconstruction on the same L1T PUPPI particles, for R=0.4 and R=0.8
- Left column: efficiency to match each anti-k_t jet to a Seeded Cone jet within $\Delta R \leq 0.2$ and p_T within 20%
- Right column: trigger efficiency as a function of simulated jet p_T for different L1T thresholds
- Seeded Cone generally matches well to anti-k_t, with some mismatches where the SC jet seeding can miss some particles / sub-jet that anti-k_t captures
- Trigger turn-ons are virtually identical for SC4 and AK4
    - Trigger rate also matches (in backup)

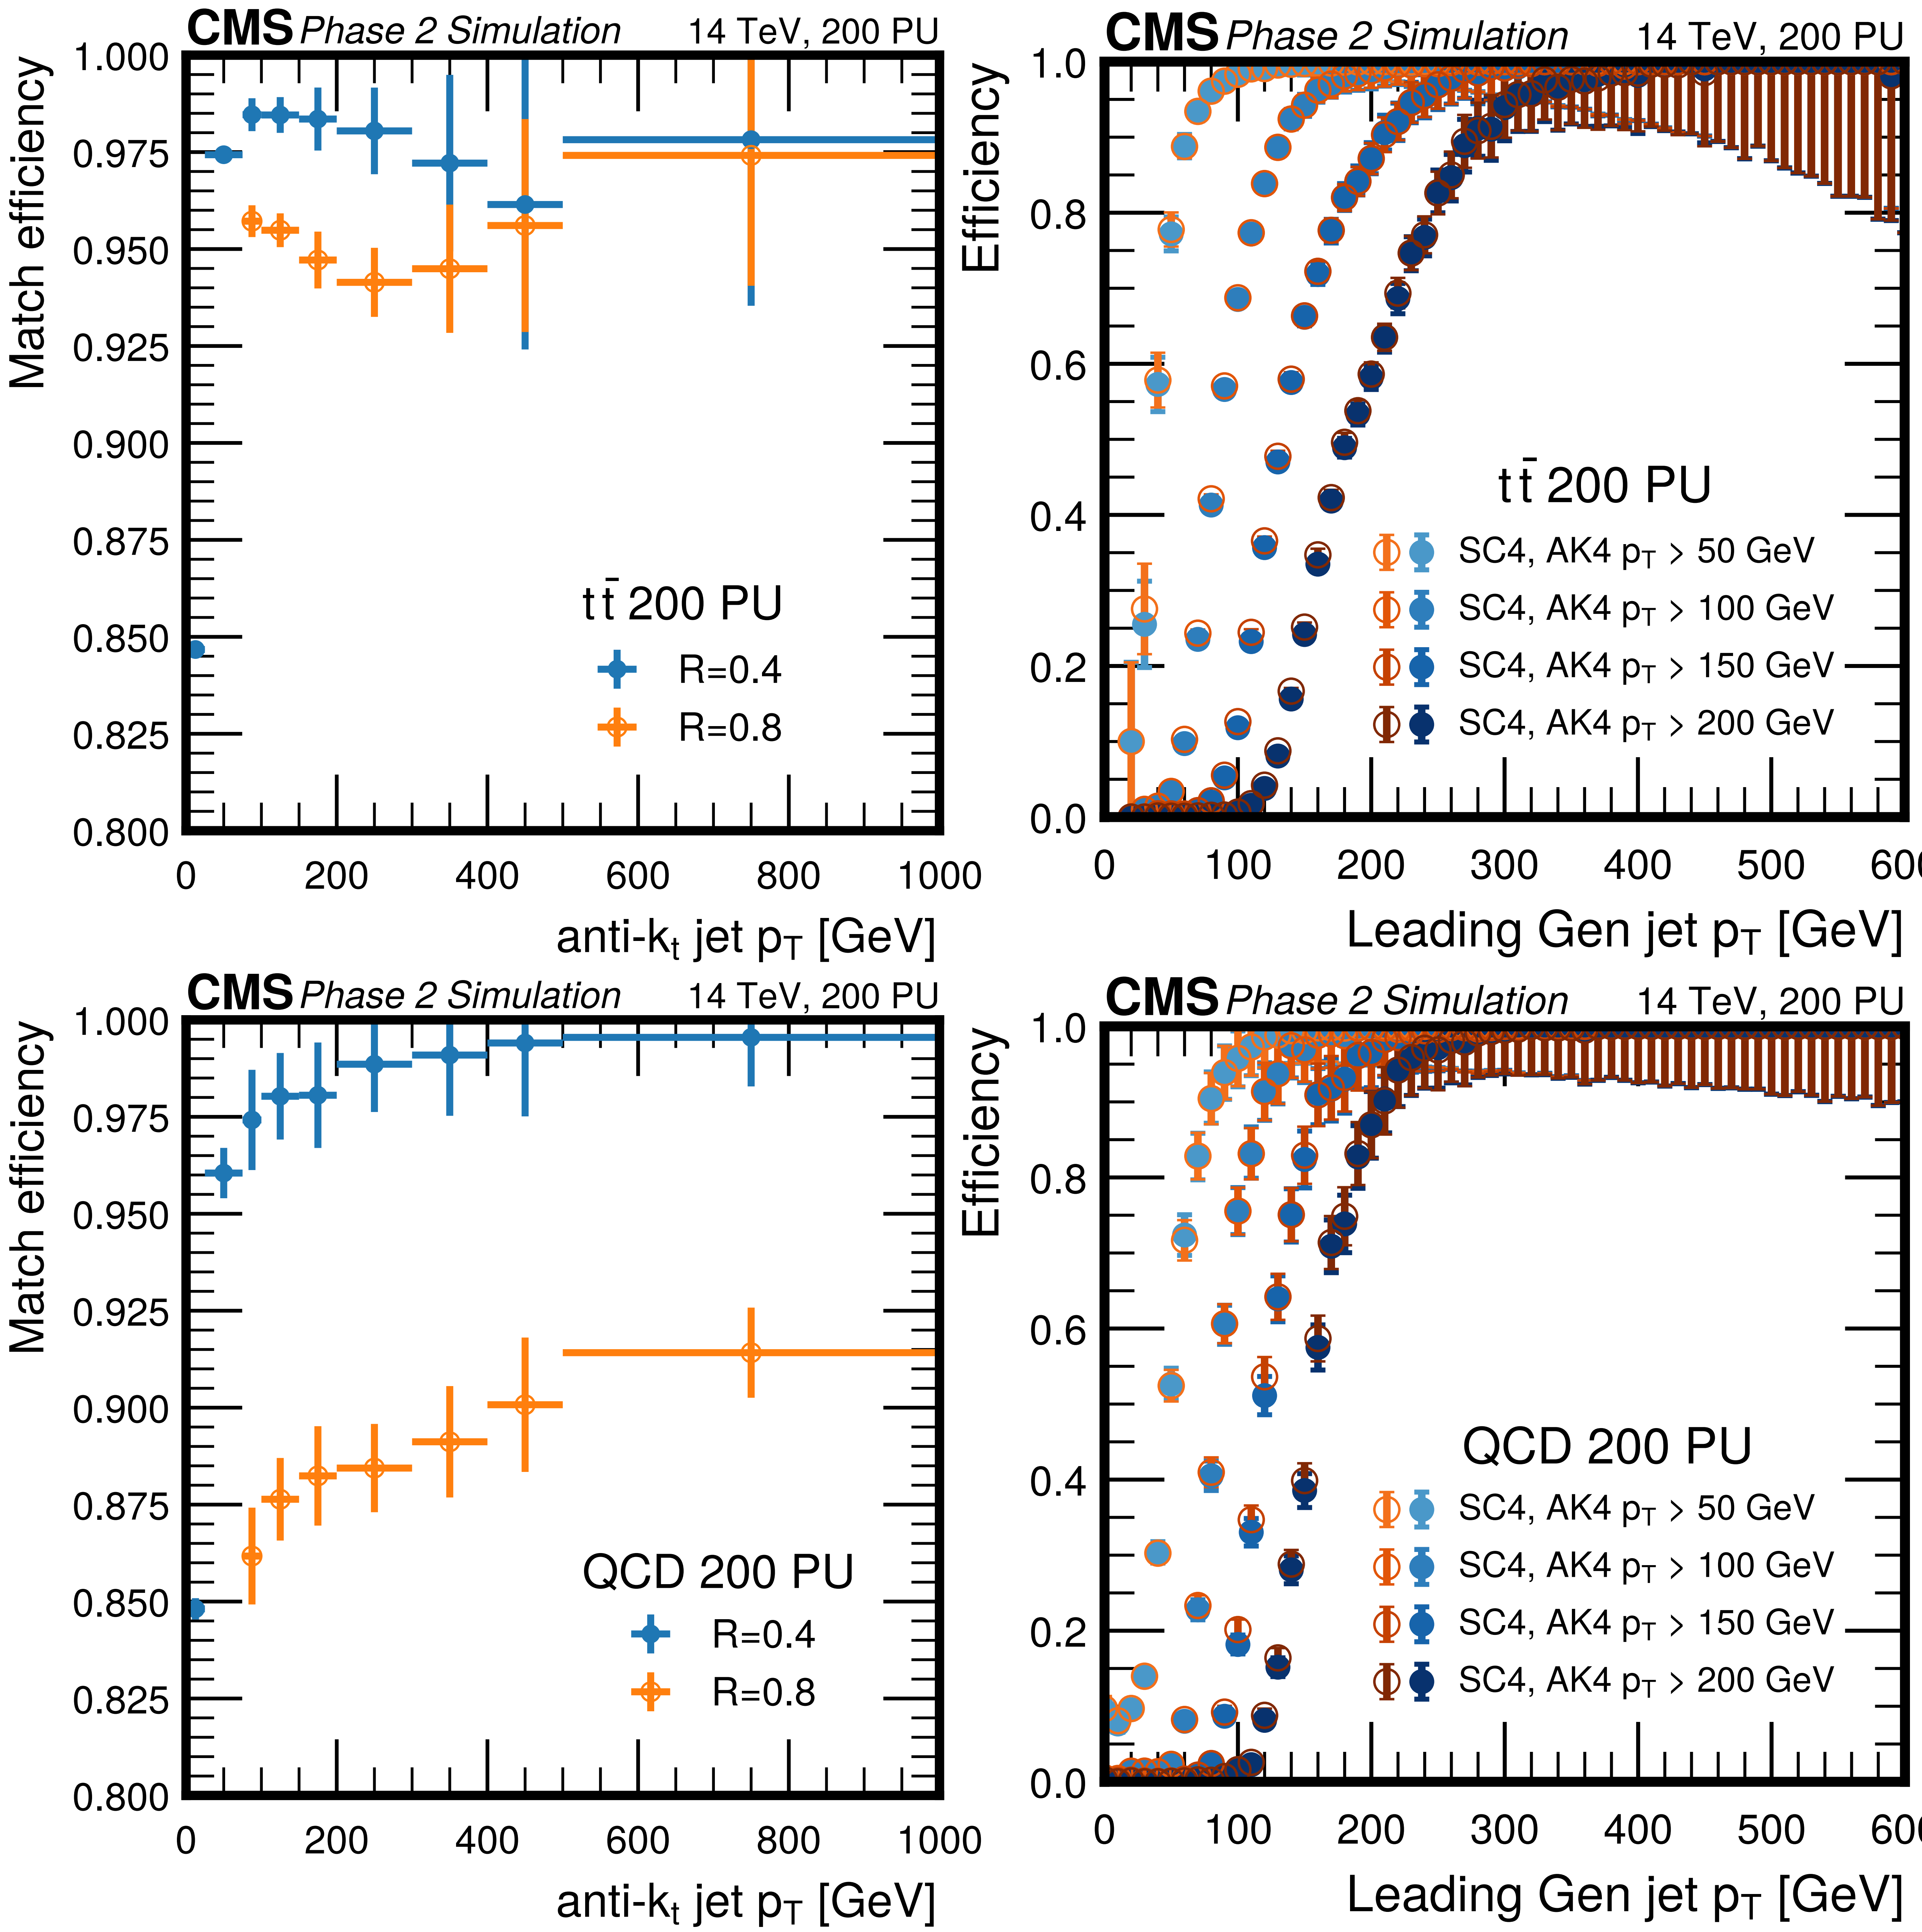

Conclusions

- The Seeded Cone jet algorithm is a lightweight, fast jet reconstruction developed for the CMS Phase-2 Level 1 Trigger
- Particles in the Correlator system are first ‘reshaped’ using a Deregionizer module for easier access
- Jets are reconstructed from the particles by first finding a high pT particle seed, then constituents within a cone of the seed
- Given L1T PUPPI particles, SC4 trigger performance is very similar to AK4
- Implementation in Xilinx VU9P FPGAs finds up to 16 jets per event with a latency of 744 ns, and processes a new event every 150 ns
    - More than 100 million jets per second per FPGA
- Presents a platform for further sophisticated jet reconstruction
    - Produce jets also with R=0.8 radius parameter for little hardware cost
    - Use jet constituent particles to perform jet tagging

Backup

Deregionizer

- Truncation of 128 particles in deregionizer motivated by multiplicity observed in high pileup simulations
- Typical event with no hard interaction and only pileup well below truncation limit
- High multiplicity topology tt with 200 pileup interactions has truncation of one particle for one event per thousand
- “Extreme scenario” tttt with 300 pileup has some more significant truncation, but many jets will be found regardless

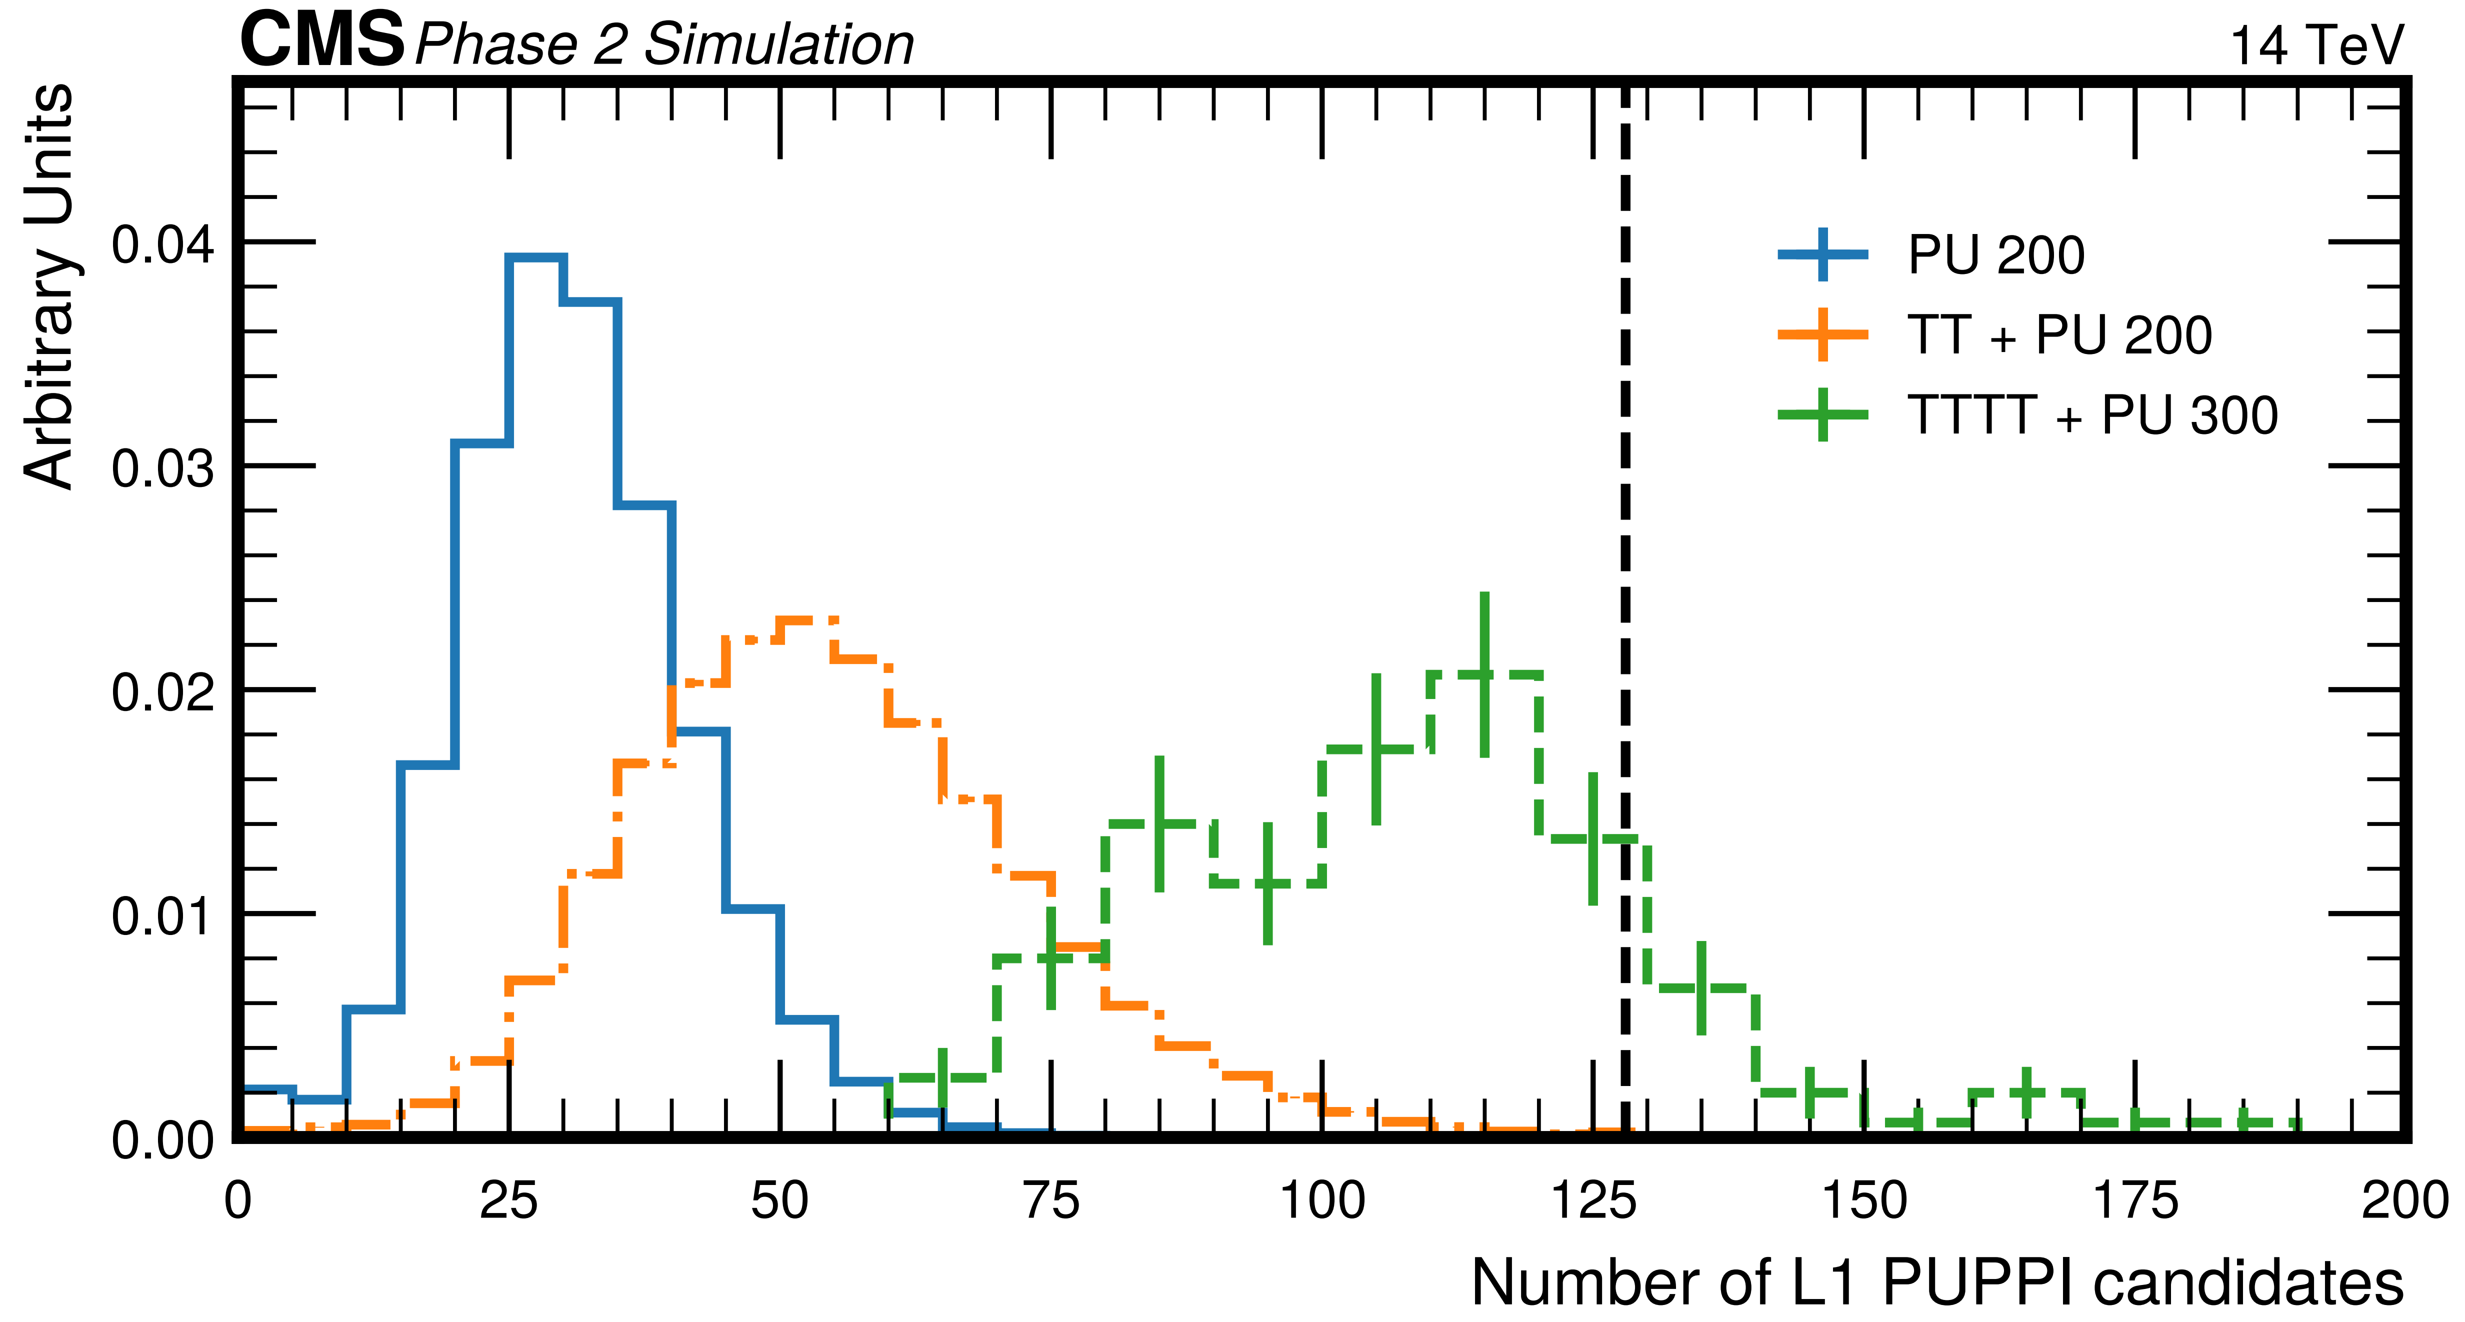

Jet performance 1 - Efficiency and Rate

- Left: online threshold vs rate in PU 200 events (no primary interaction)
- Centre: turn-on curve with thresholds chosen for a rate of 70 kHz in $t\bar{t}$ with 200 pileup
- Right: turn-on curve with thresholds chosen for a rate of 70 kHz in QCD with 200 pileup
- SC4 performance nearly identical to AK4

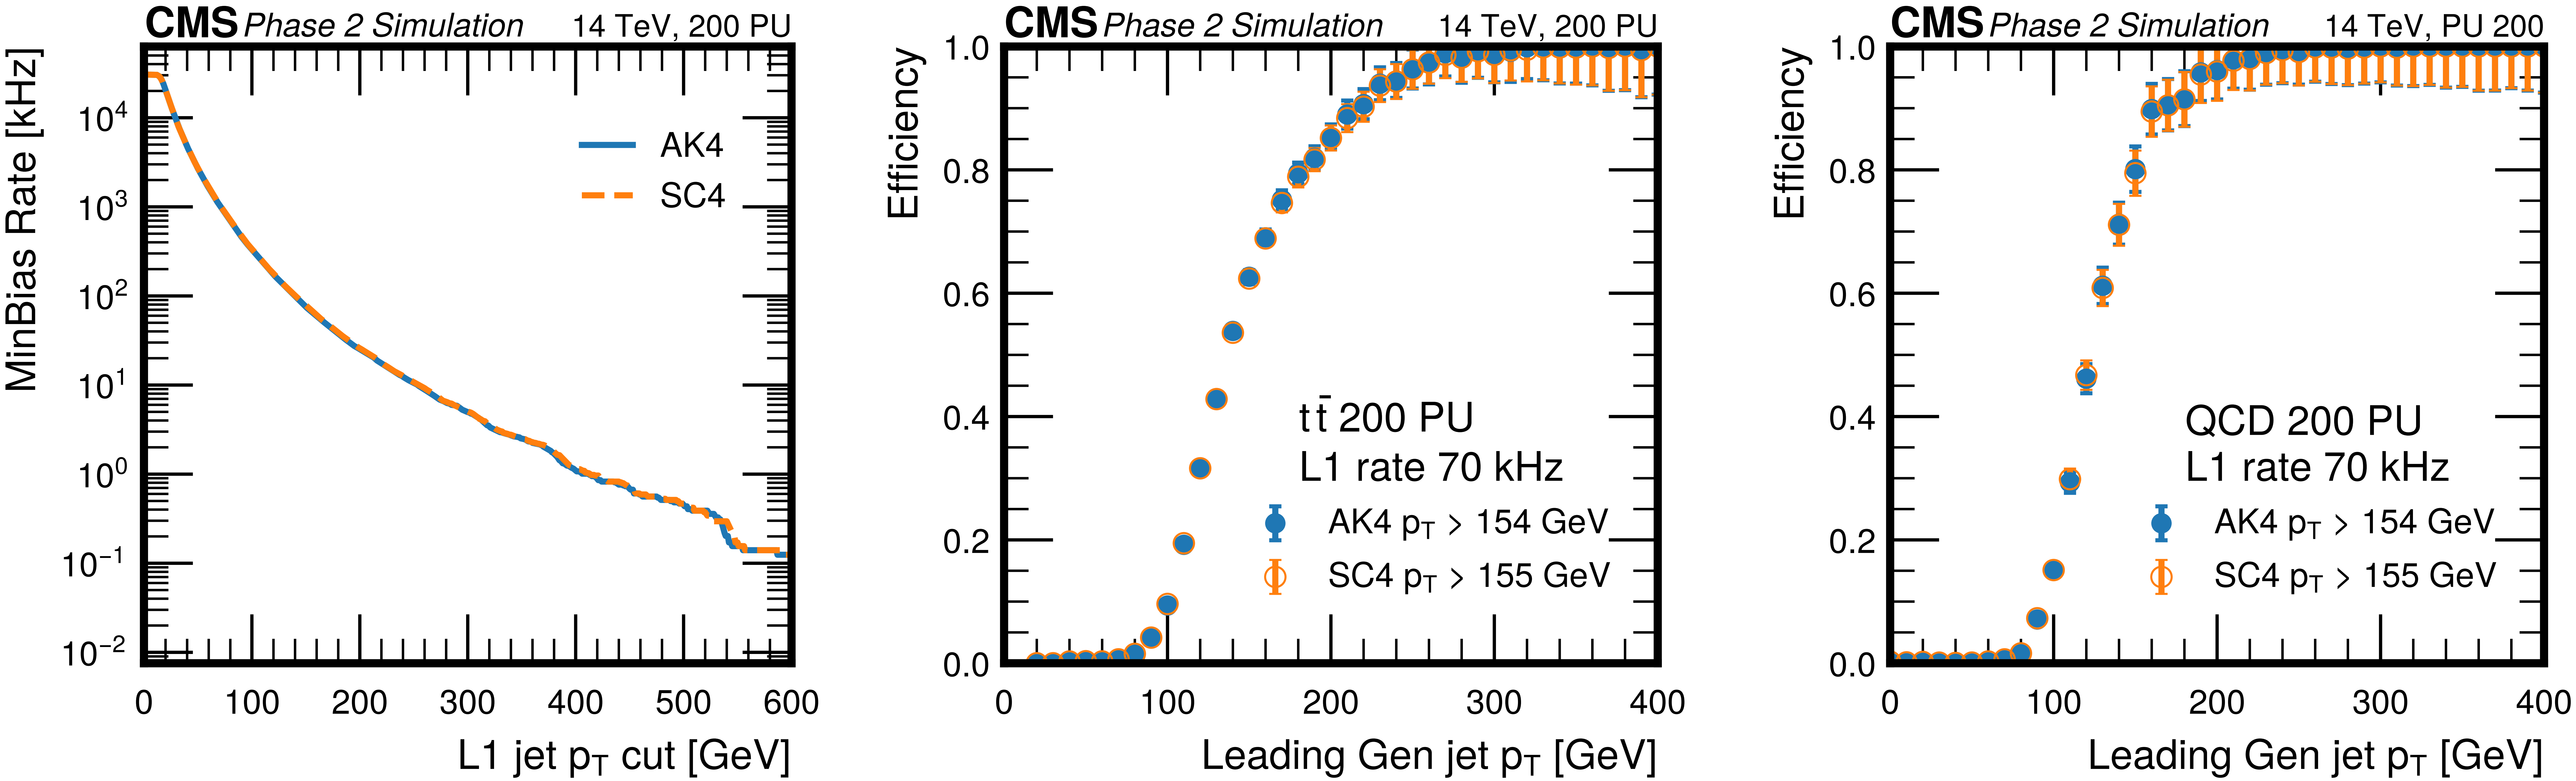

Jet performance 4 - SC to AK matching

Left: distribution of Seeded Cone p_T for jets matched within $\Delta R \leq 0.2$ of an anti-k_t jet with R=0.4 in simulated events of $t\bar{t}$ with 200 pileup.

Right: distribution of Seeded Cone p_T for jets matched within $\Delta R \leq 0.2$ of an anti-k_t jet with R=0.8 in simulated events of $t\bar{t}$ with 200 pileup.

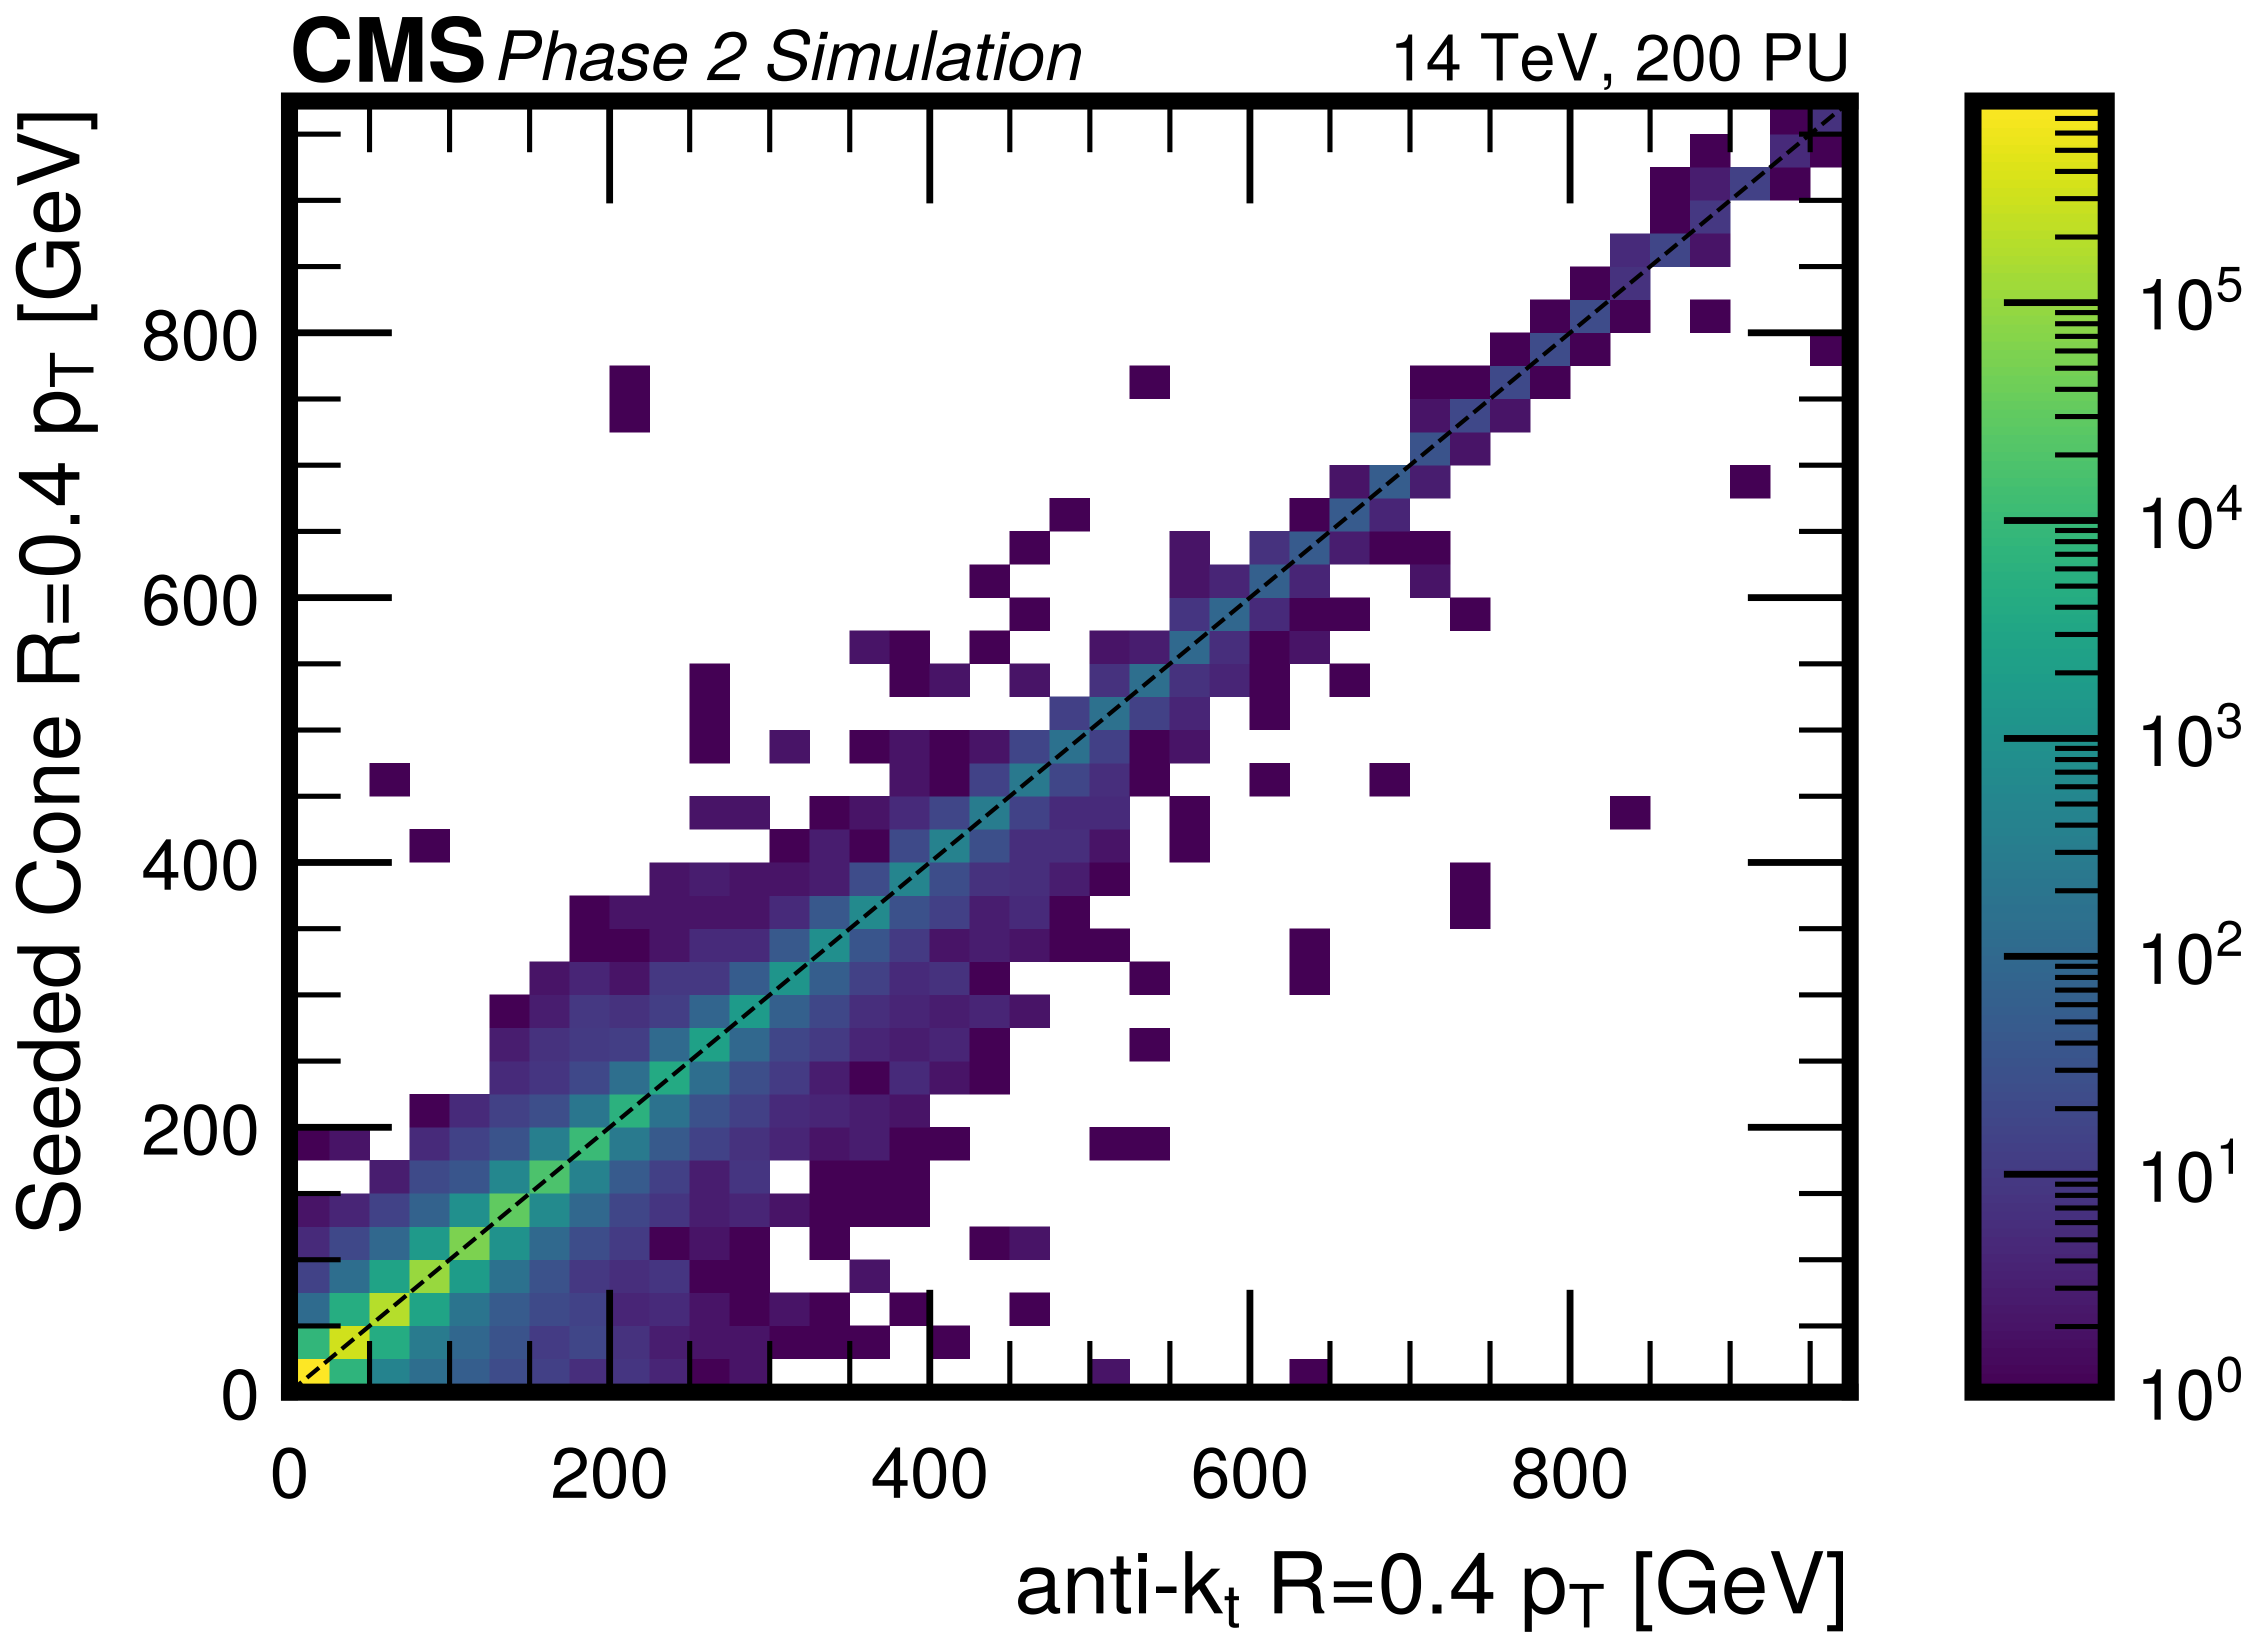

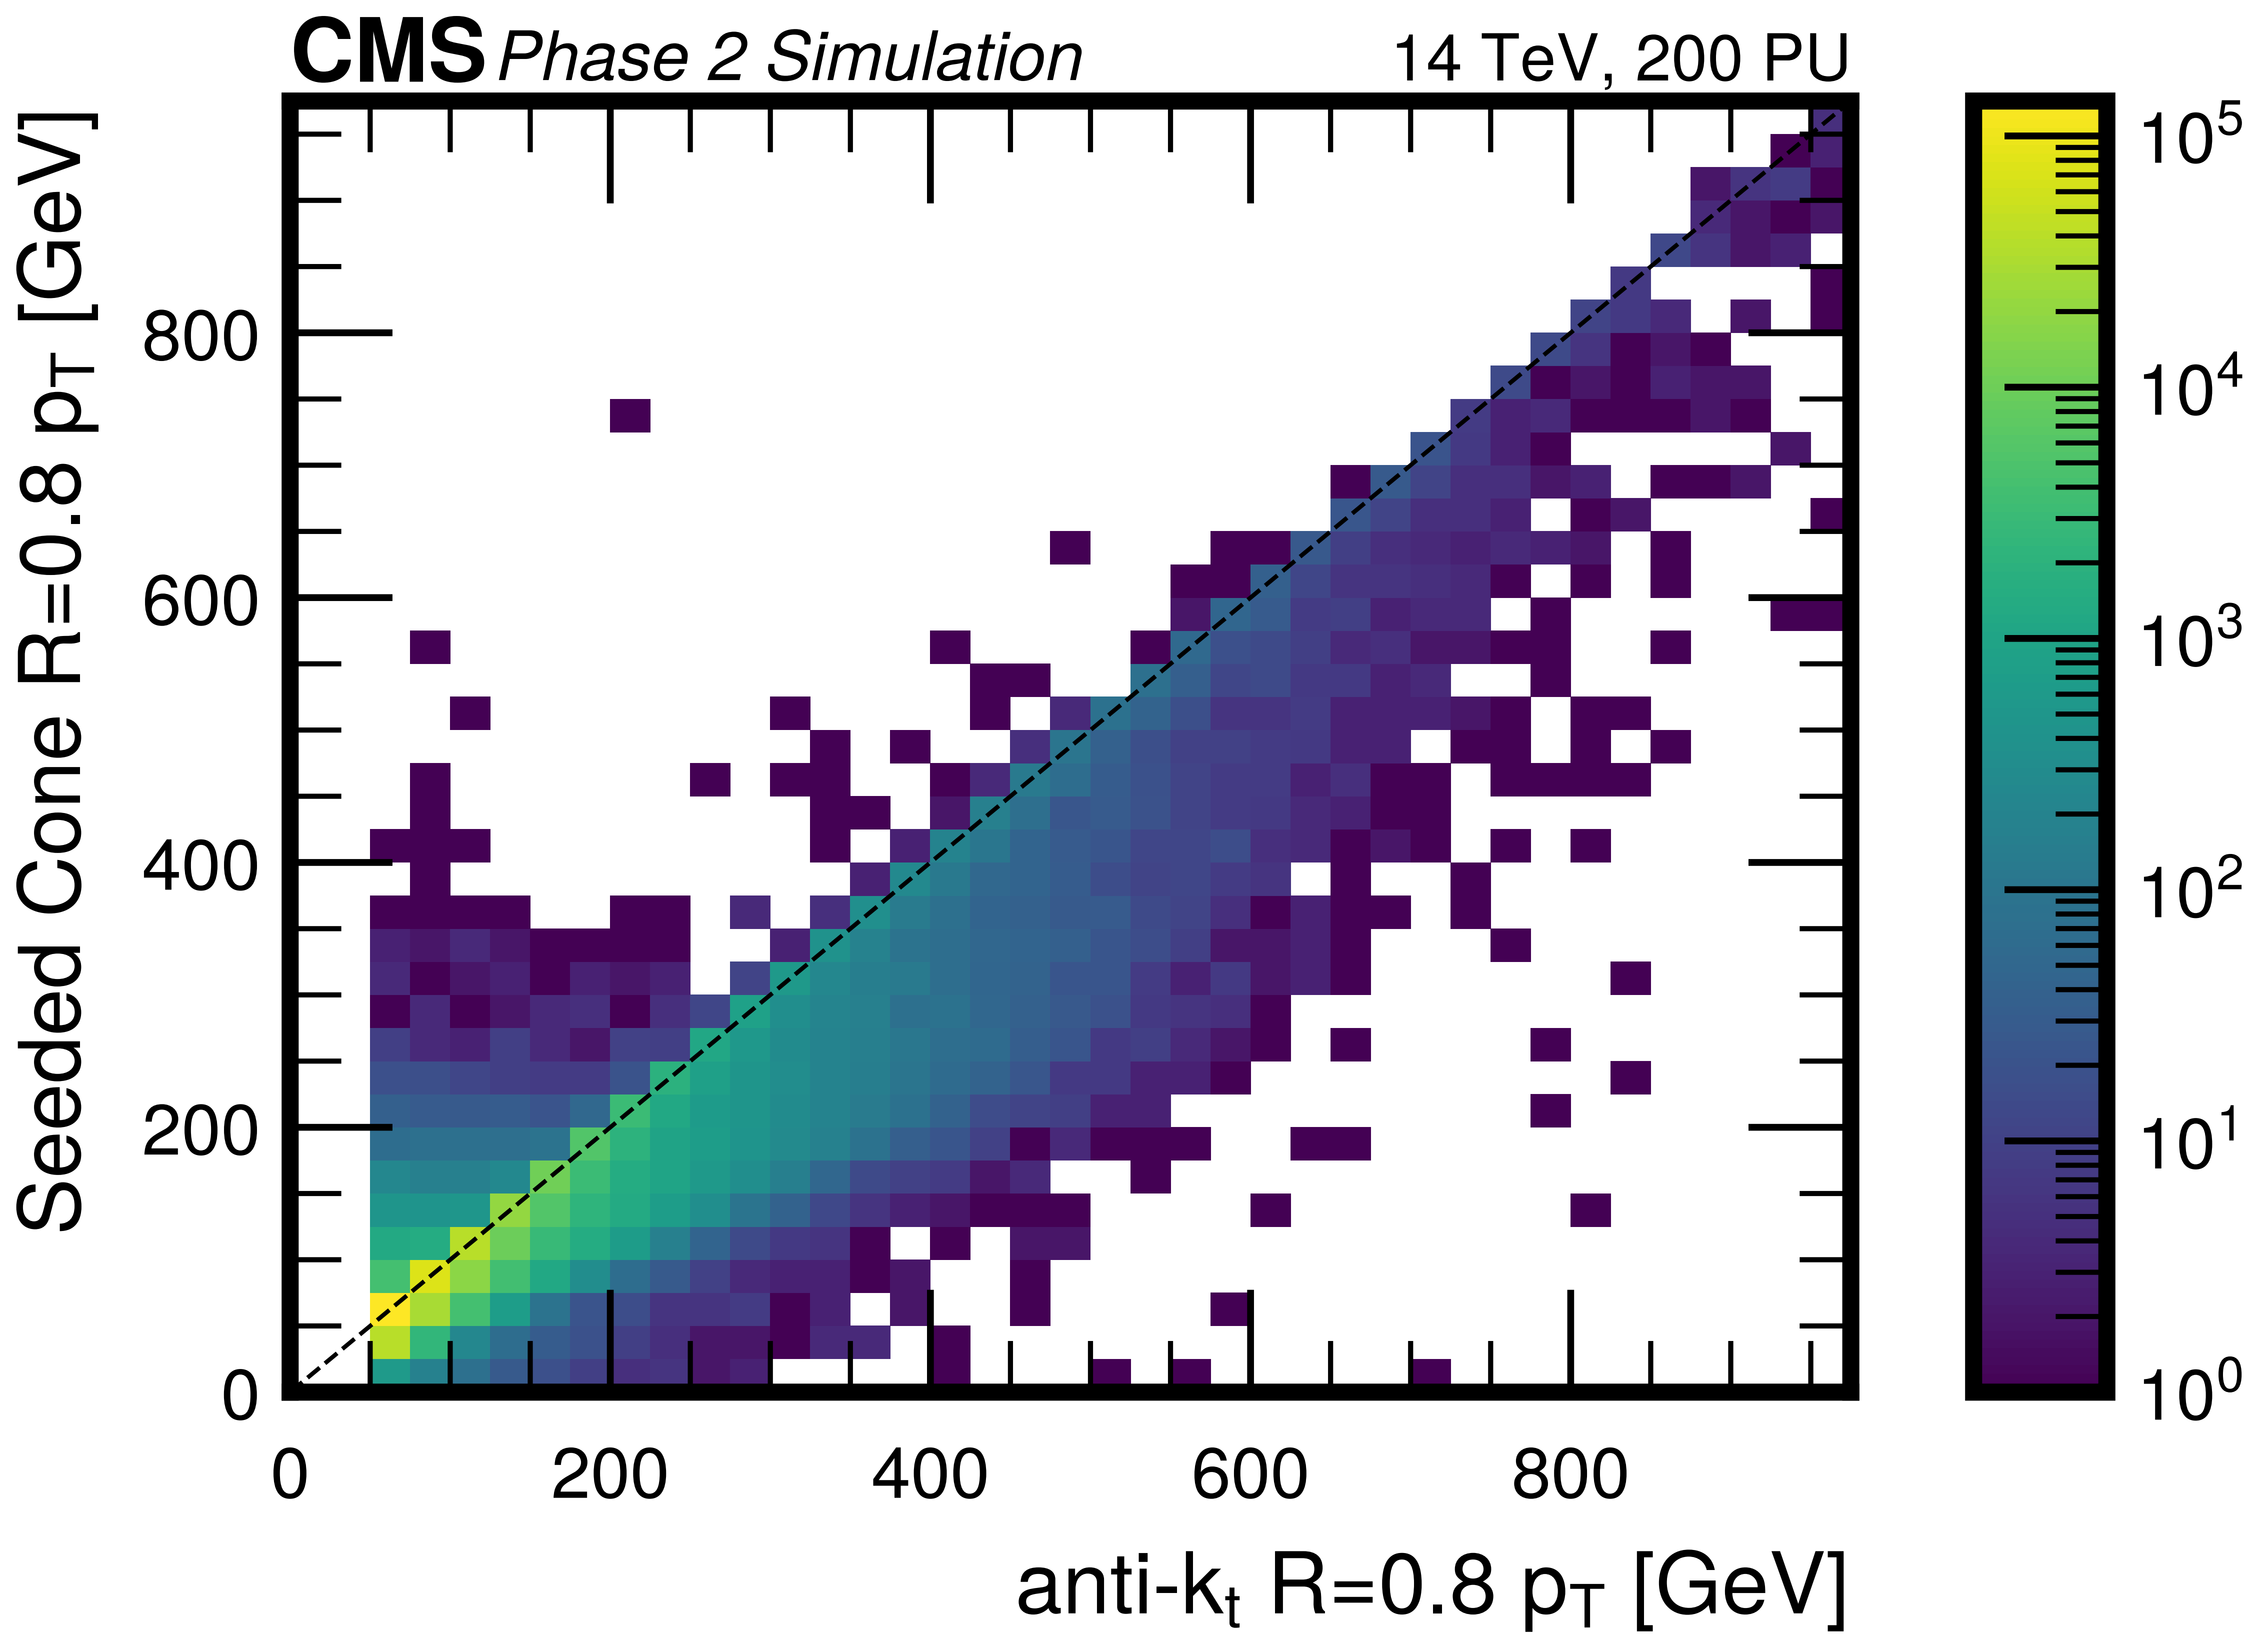

Event Displays 1

- “Event display” from sample of Higgs bosons produced in association with a W or Z boson
- Shows a Higgs with p_T of 400 GeV decaying to two b quarks around $(\eta,\phi) \simeq (0,1)$ and a W boson with p_T of 360 GeV decaying to light quarks around $(\eta,\phi) \simeq (-1,-2)$
- The W is reconstructed as two R=0.4 jets or one R=0.8 jet by both Seeded Cone and anti-k_t reconstruction.

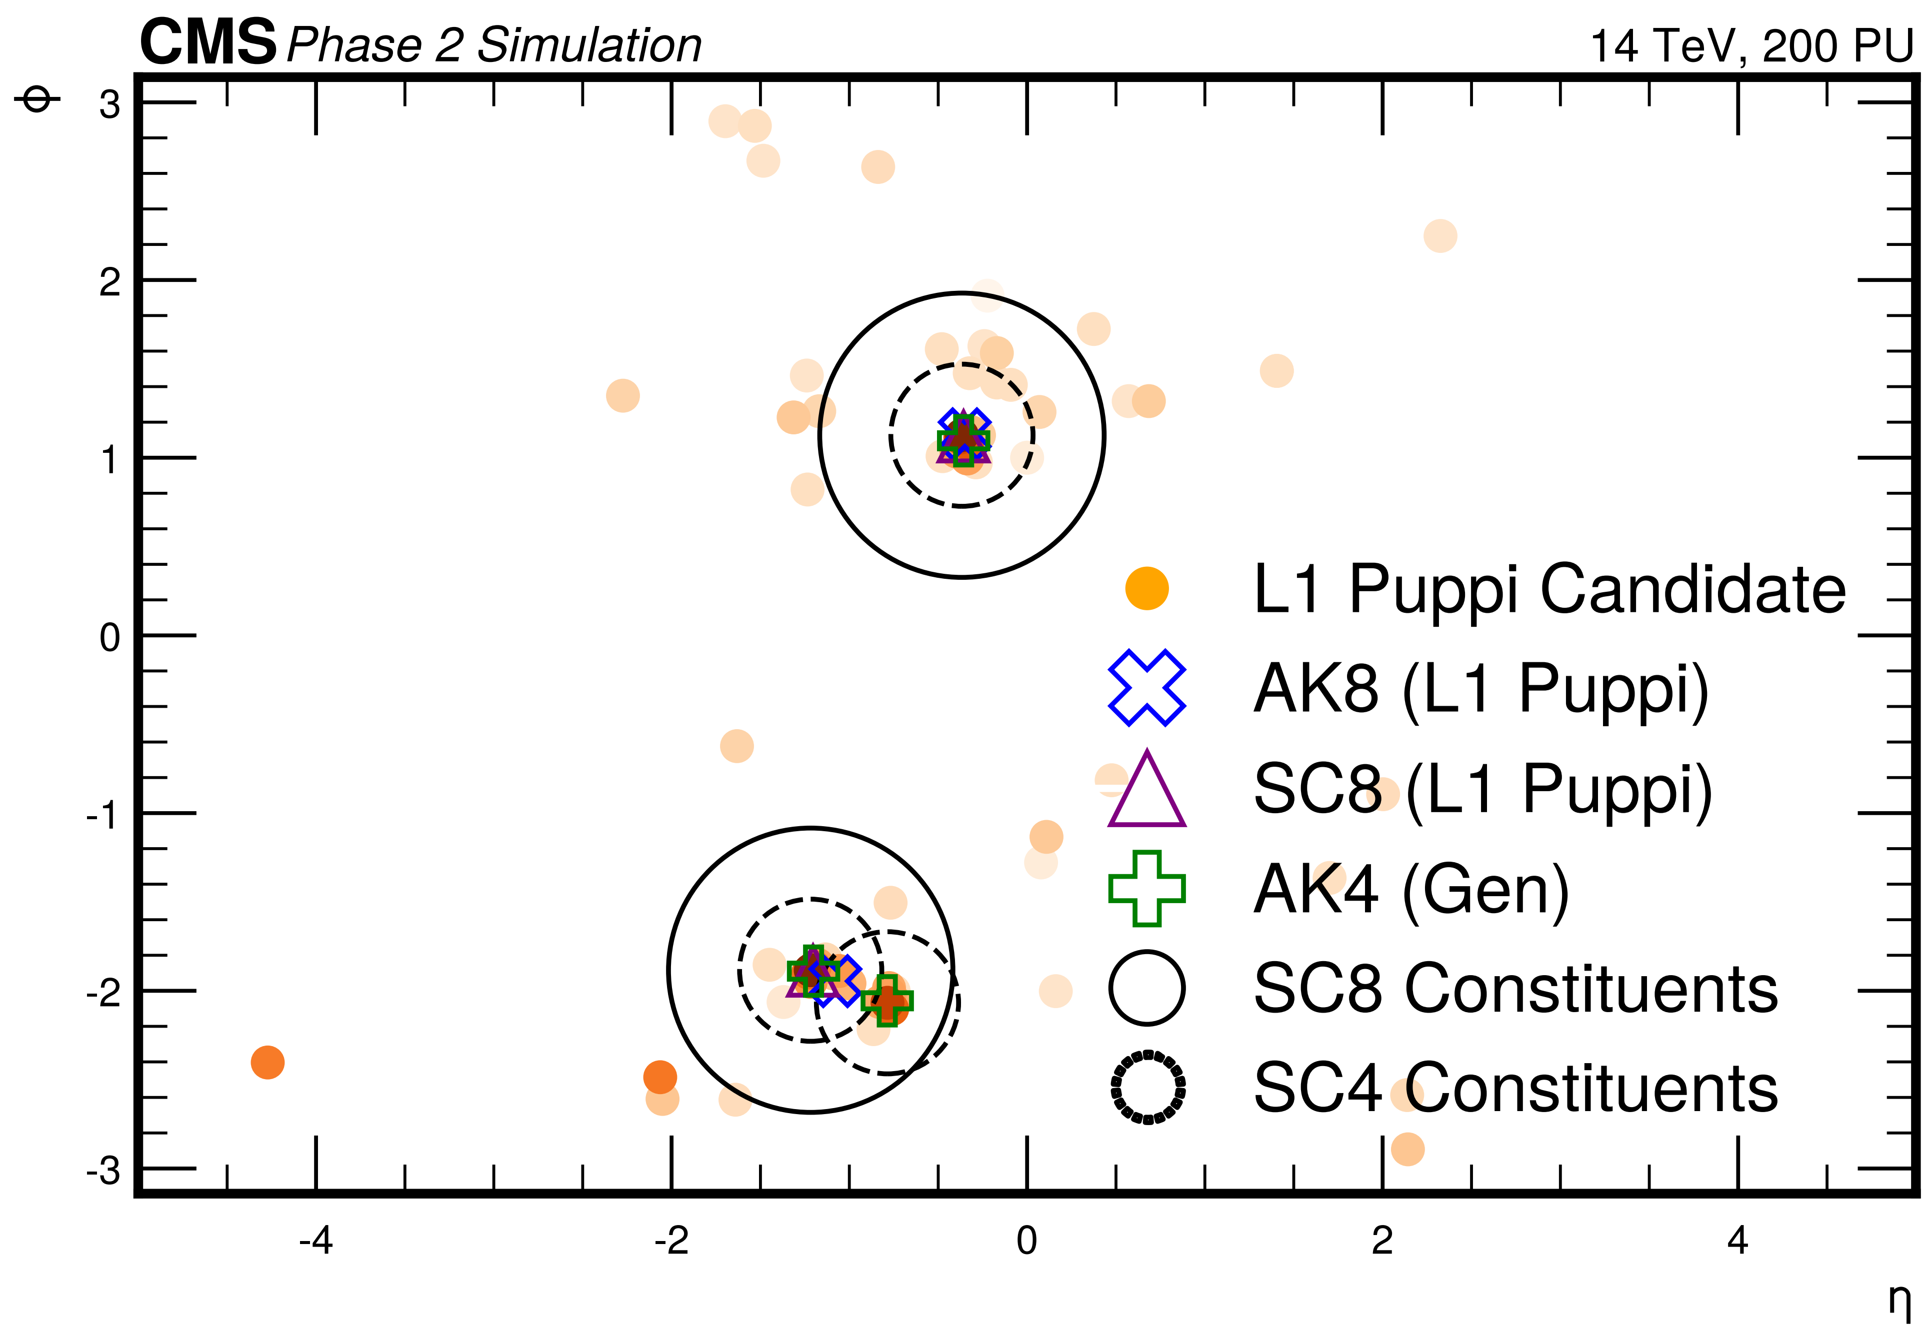

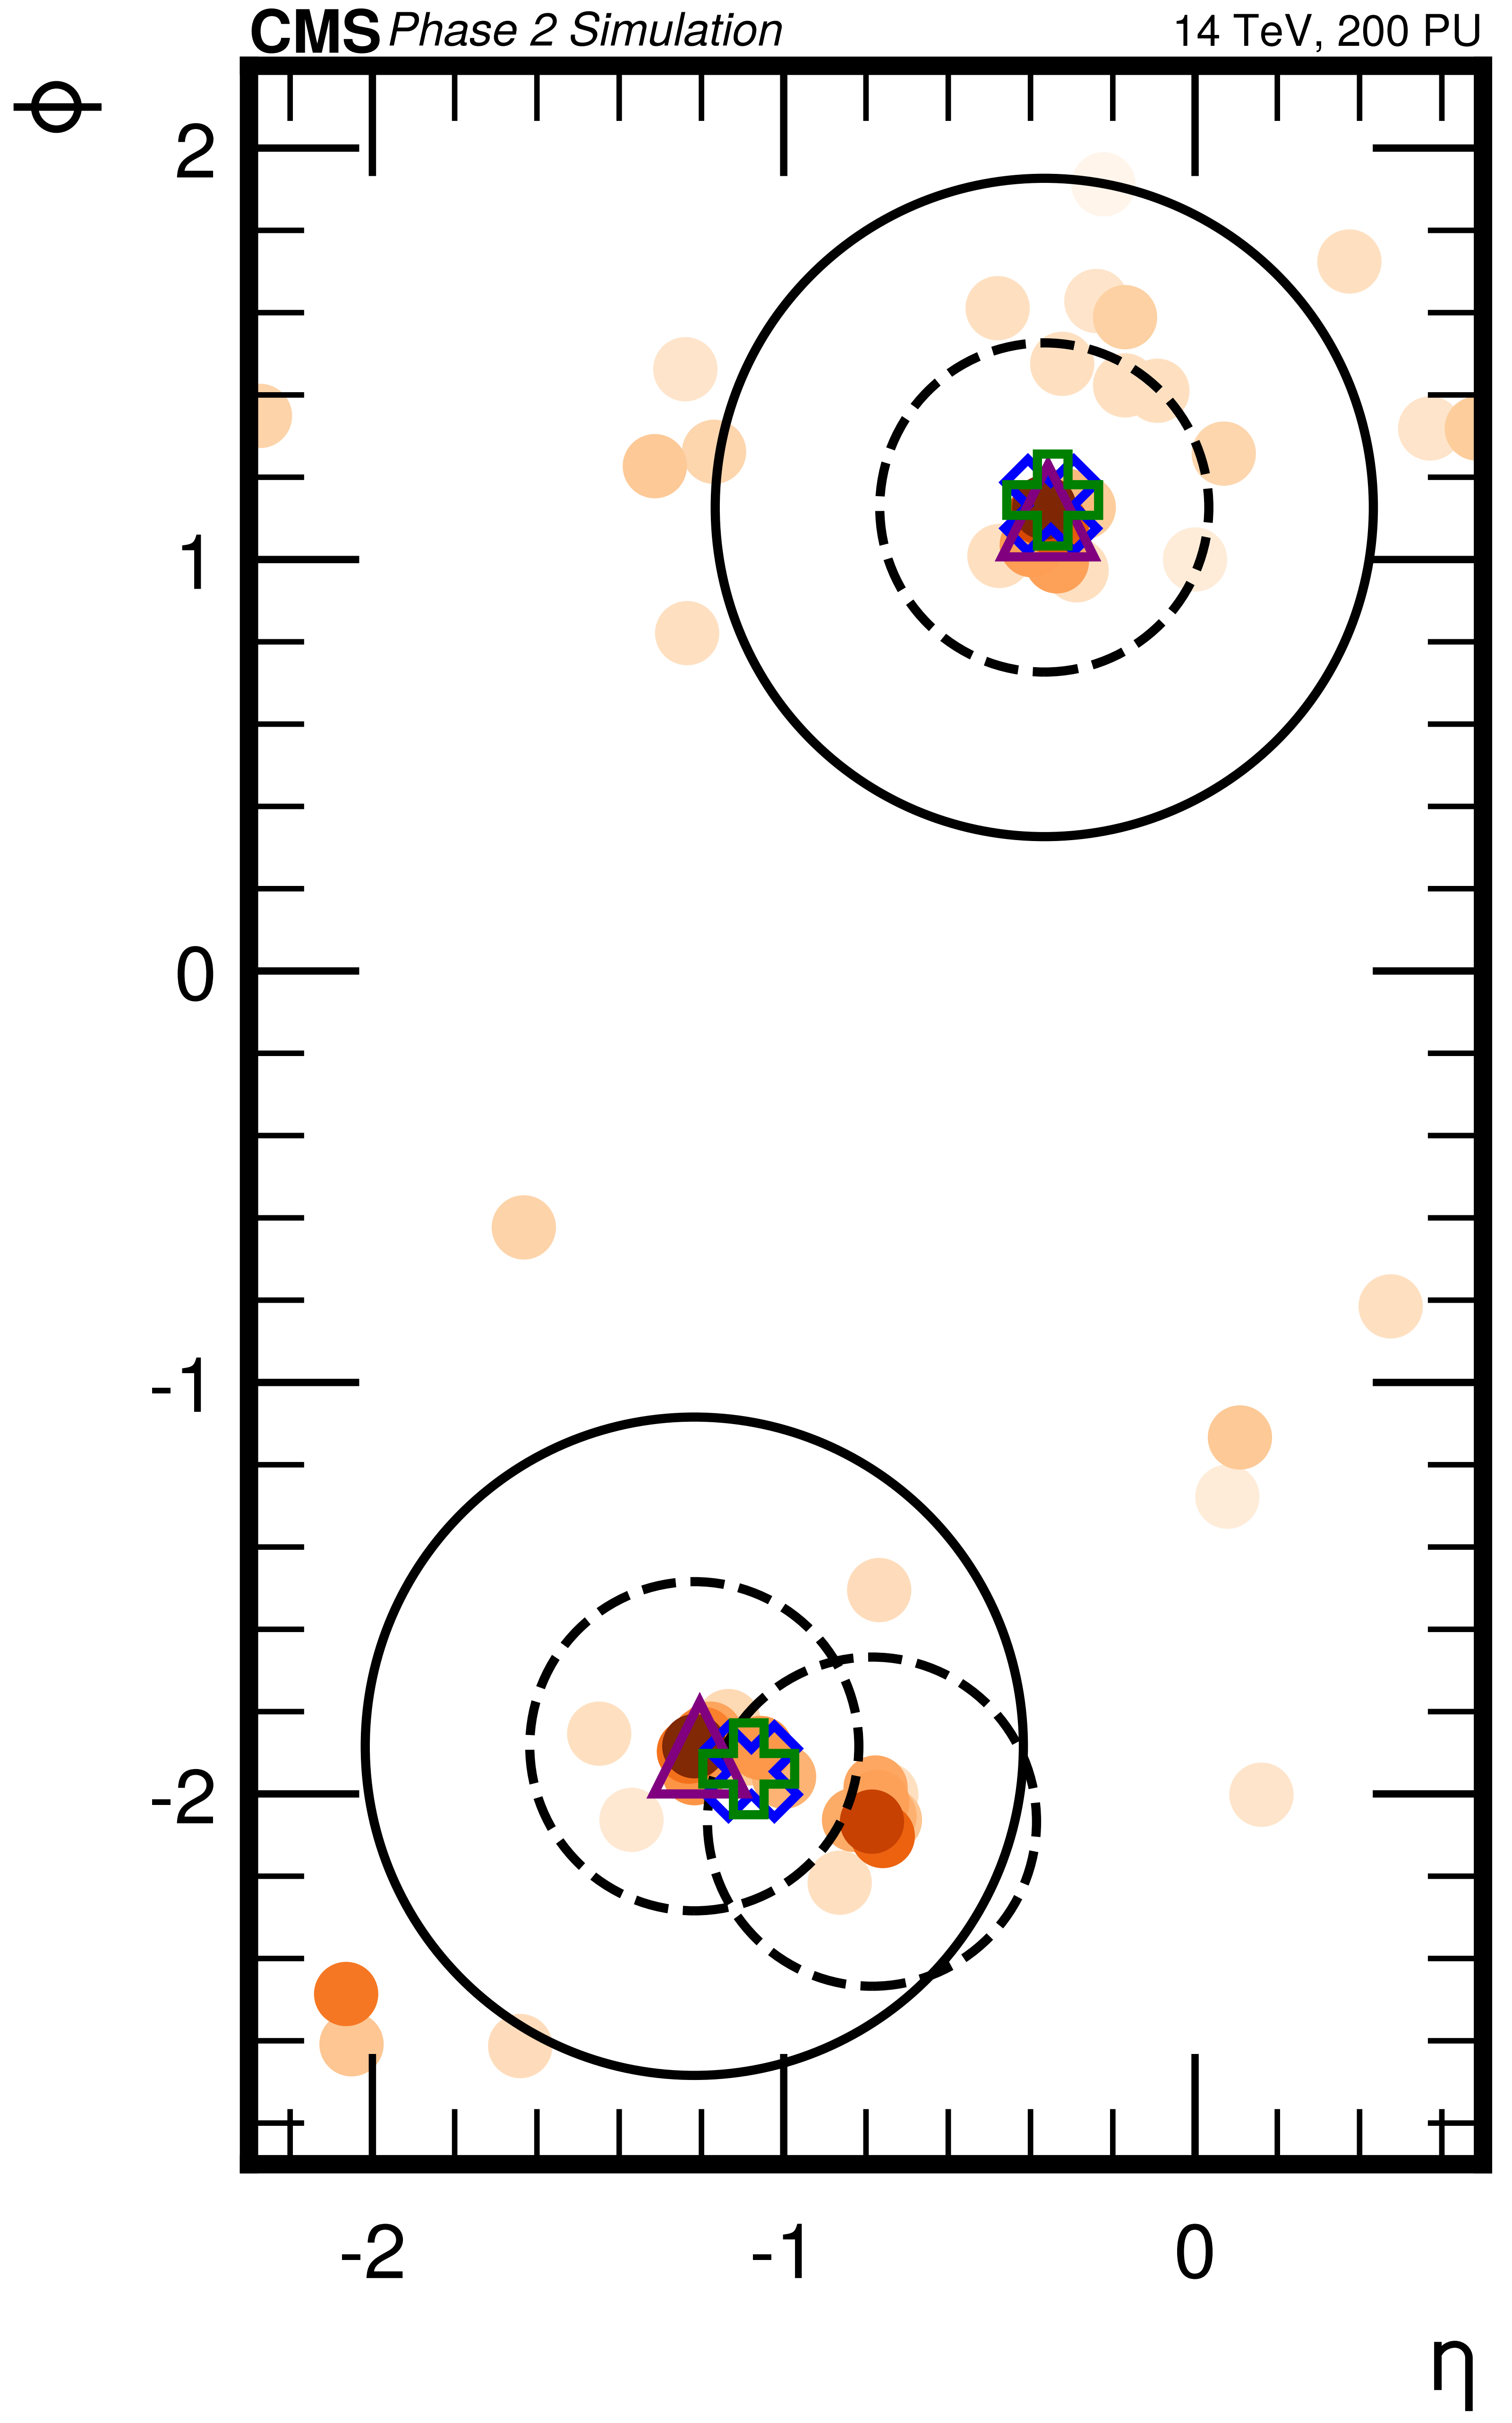

Event Displays 2

- “Event display” from sample of $t\bar{t}$ with 200 pileup
- Shows a case where one AK8 jet is reconstructed as two SC8 jets due to the limitations of casting the cone around the single highest p_T particle seed with no reclustering

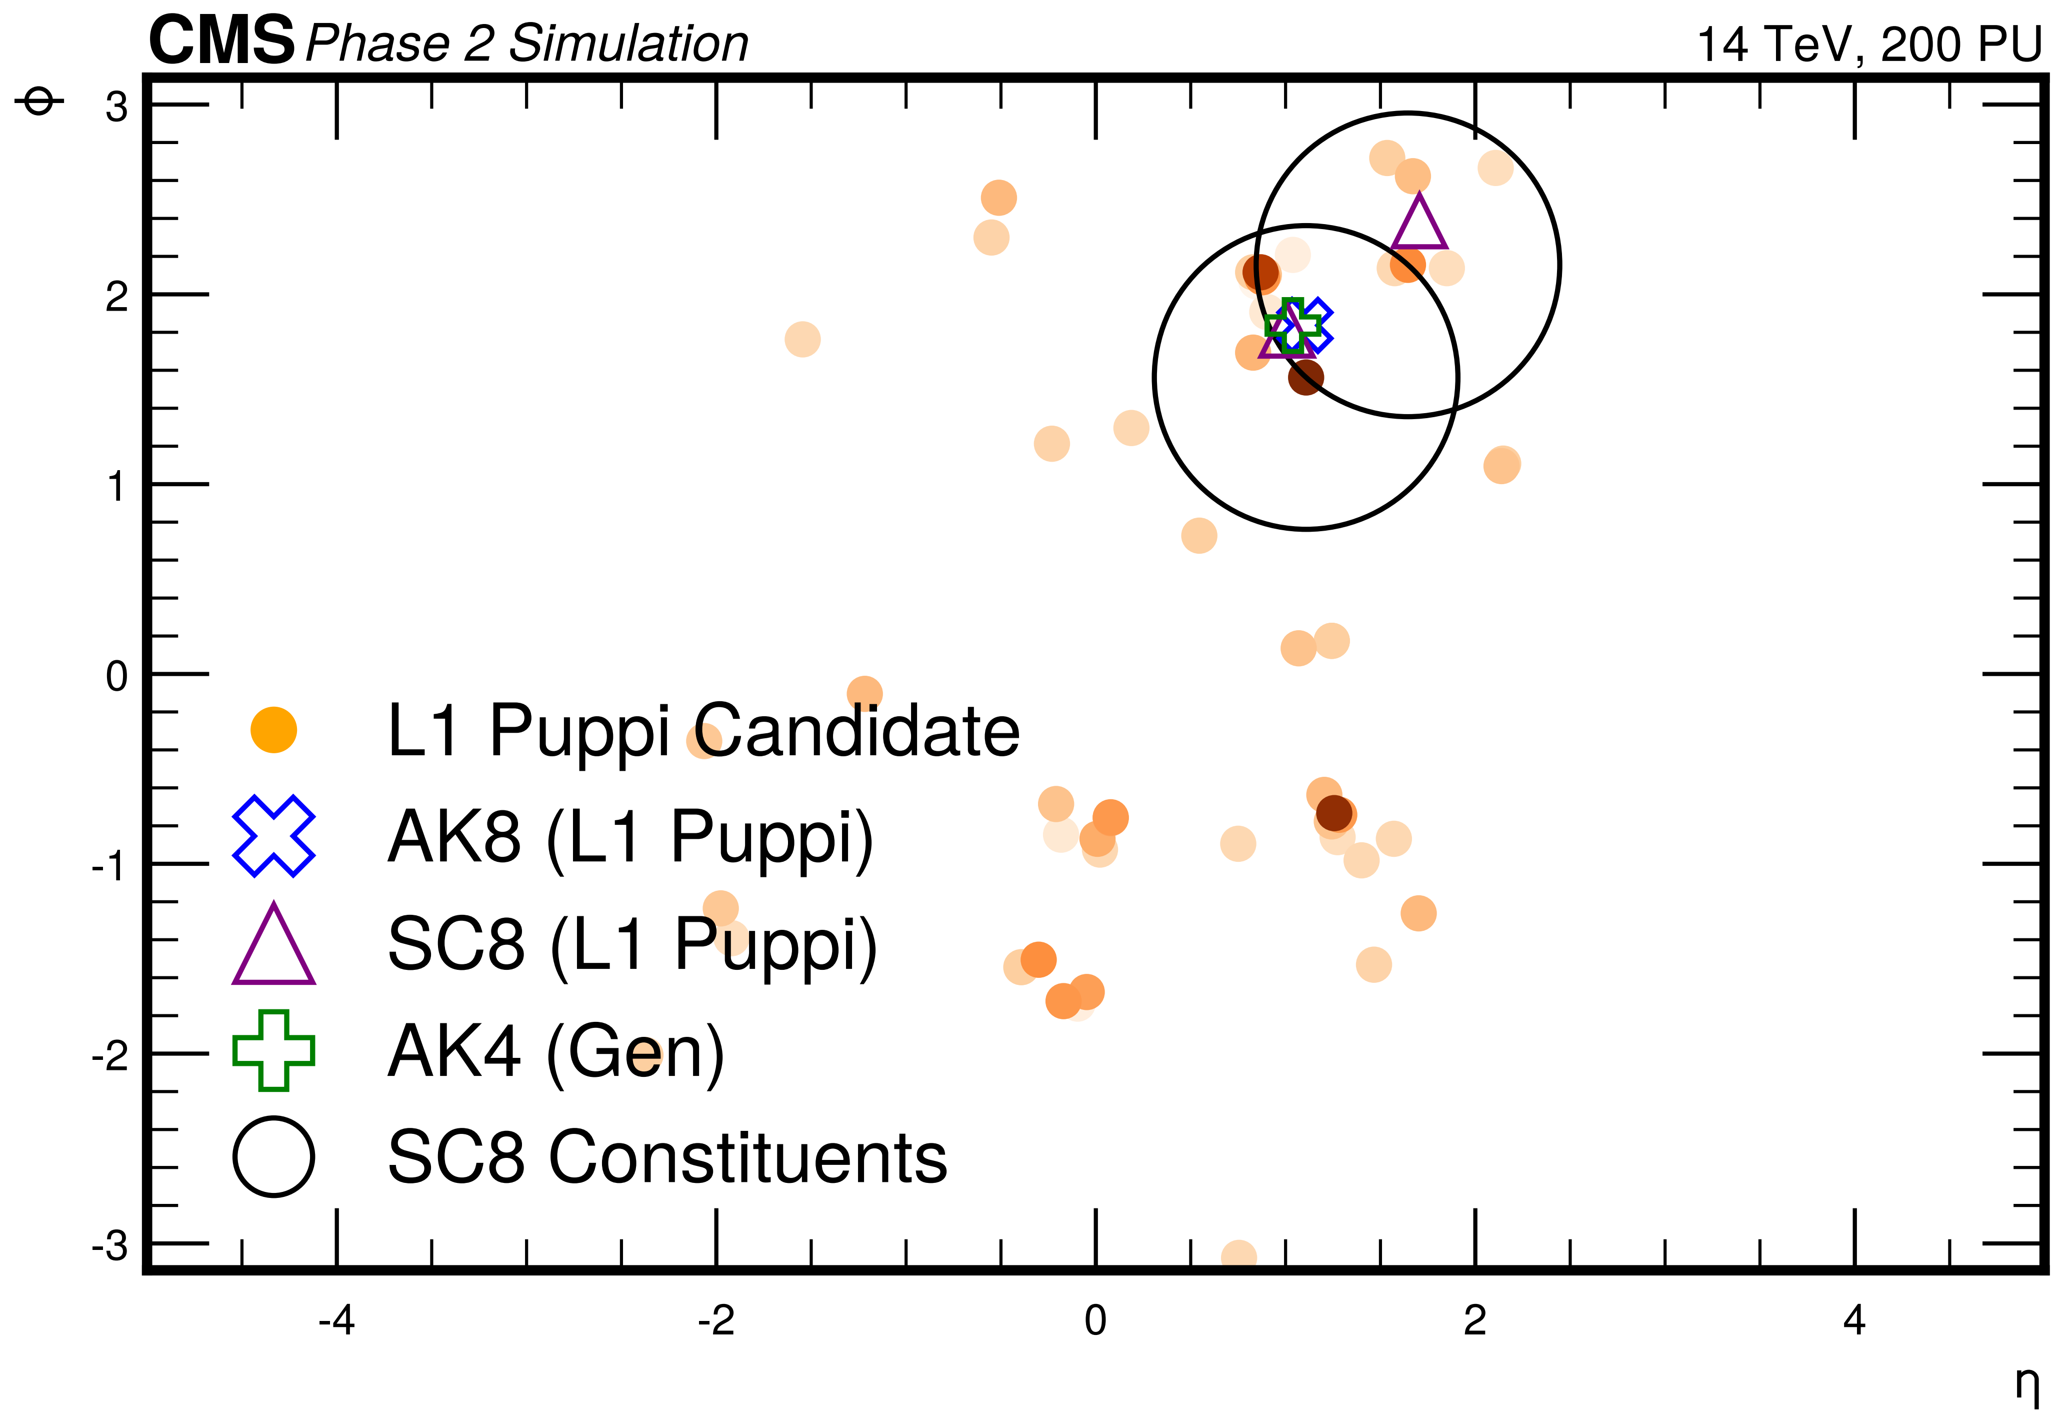

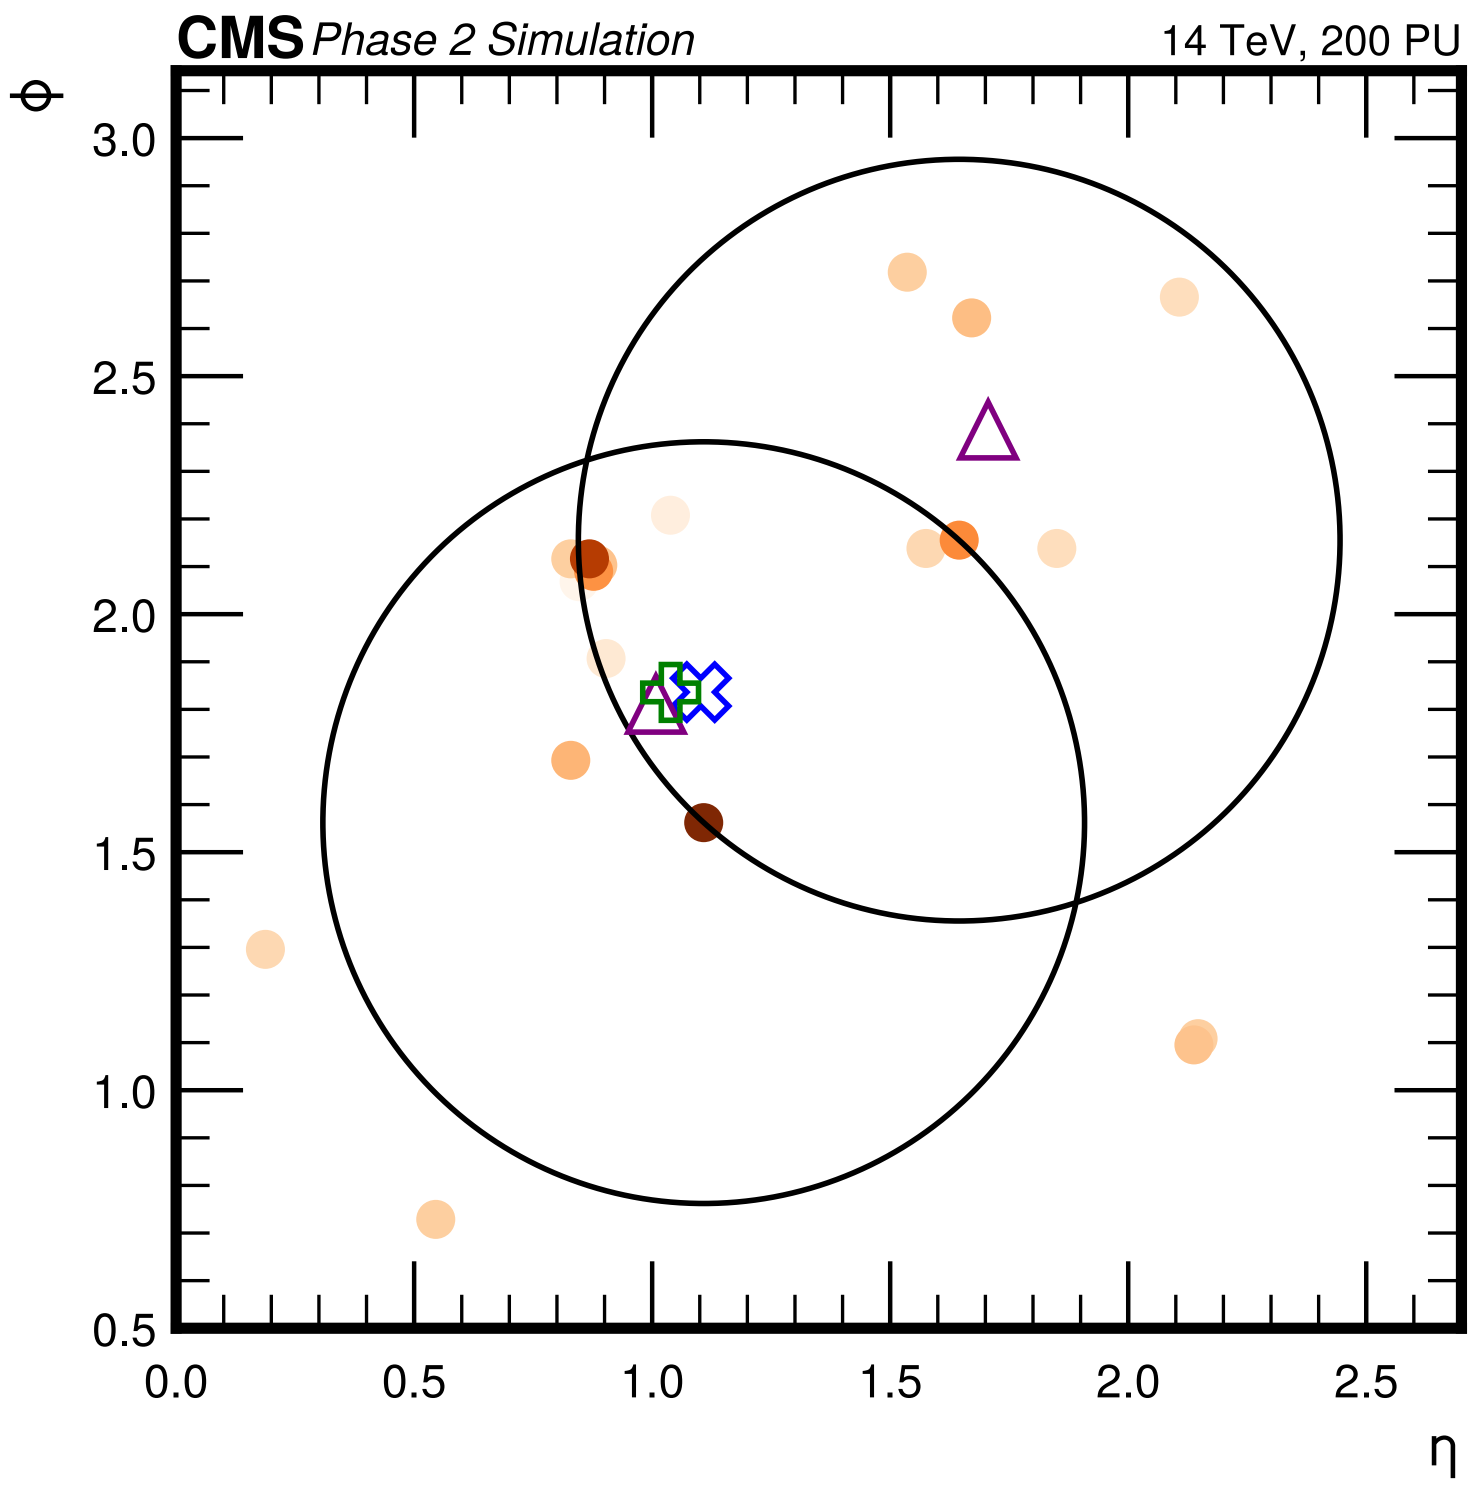

Event Displays

- Event displays from $t\bar{t}$ with 200 PU where one Gen AK4 jet is resolved as two SC4/AK4 jets or one SC8/AK8 jet