



# Updates

*Viktor Khristenko (CERN), Maria Girone (CERN)*

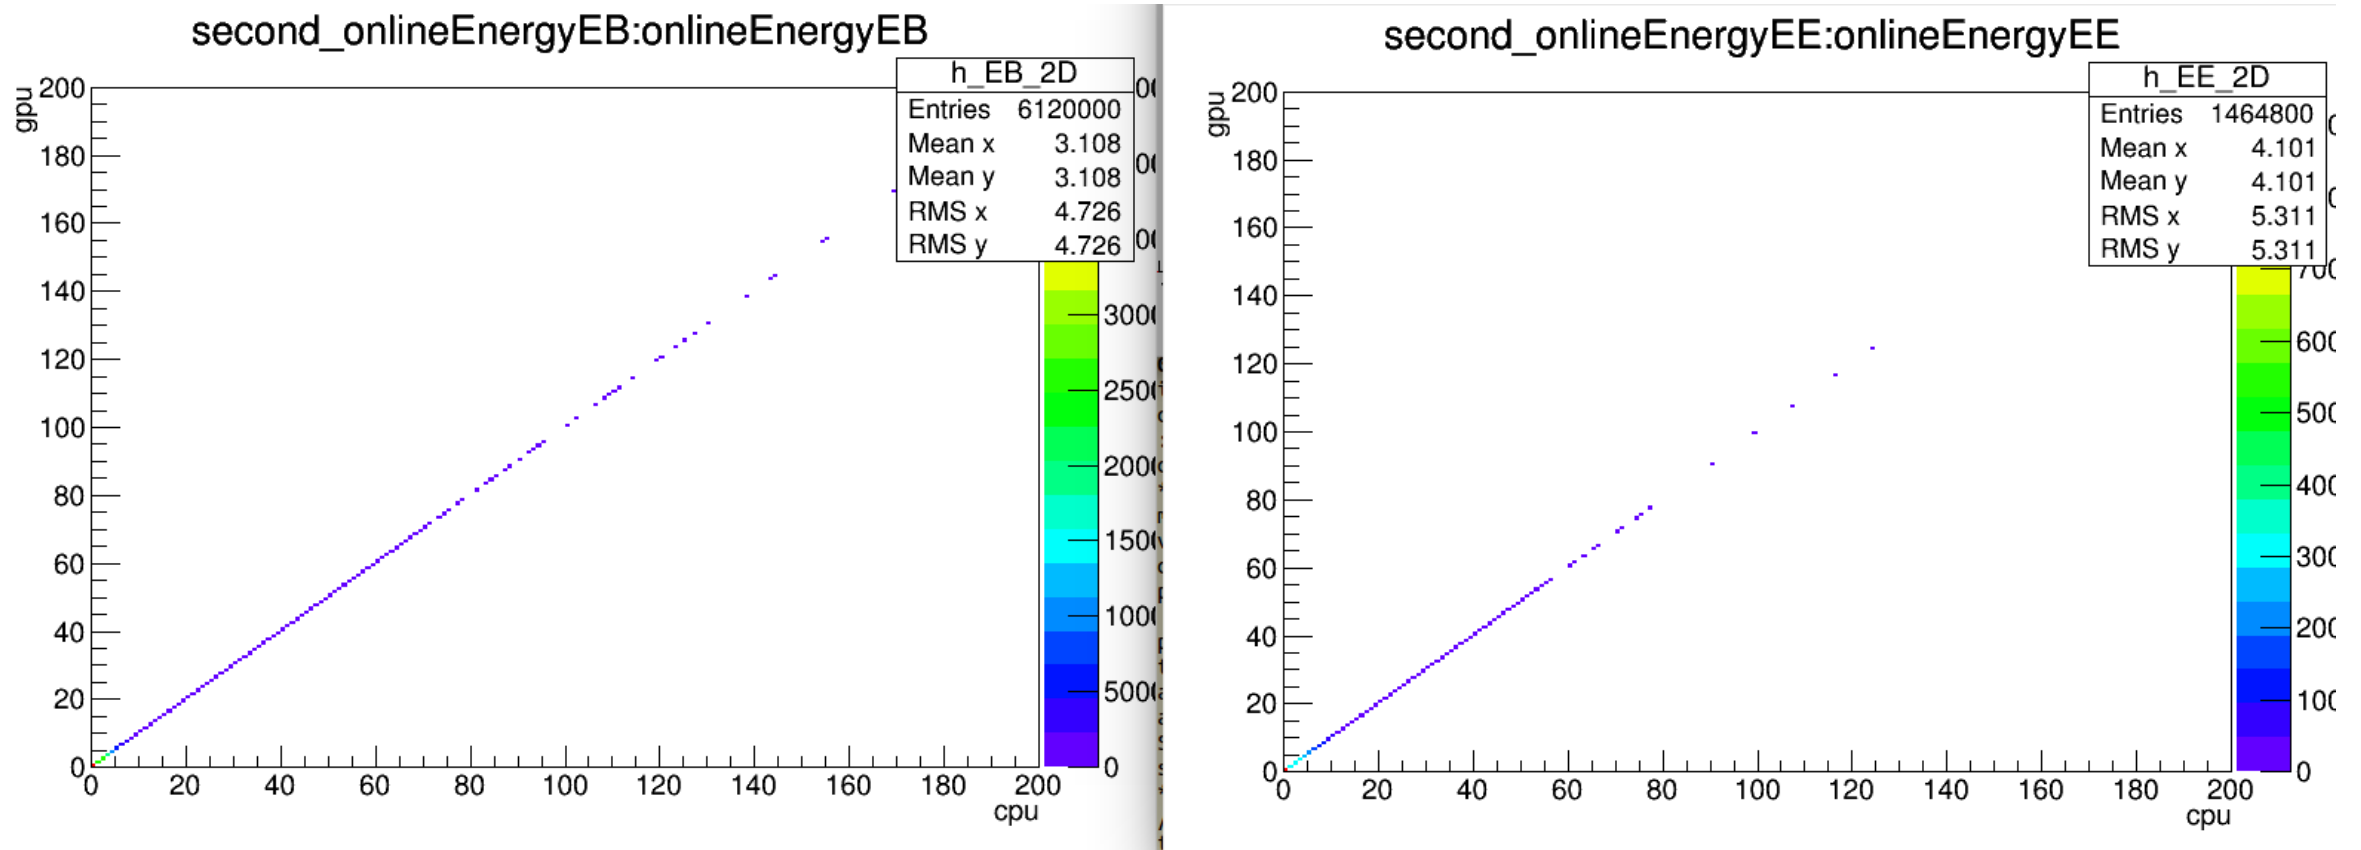


# Ecal Reco for GPU

- Just to remind, main activity here is to port existing CPU workload to utilize GPUs and optimize.
  - We are embarrassingly parallel -> no rpc/ipc
- Full Electromagnetic Calorimeter Reco with CUDA implemented
  - There are a few addons, not required now but will be needed in the upcoming future of LHC
- Integrated with CMSSW
  - Fully integrated with the experiment's framework.

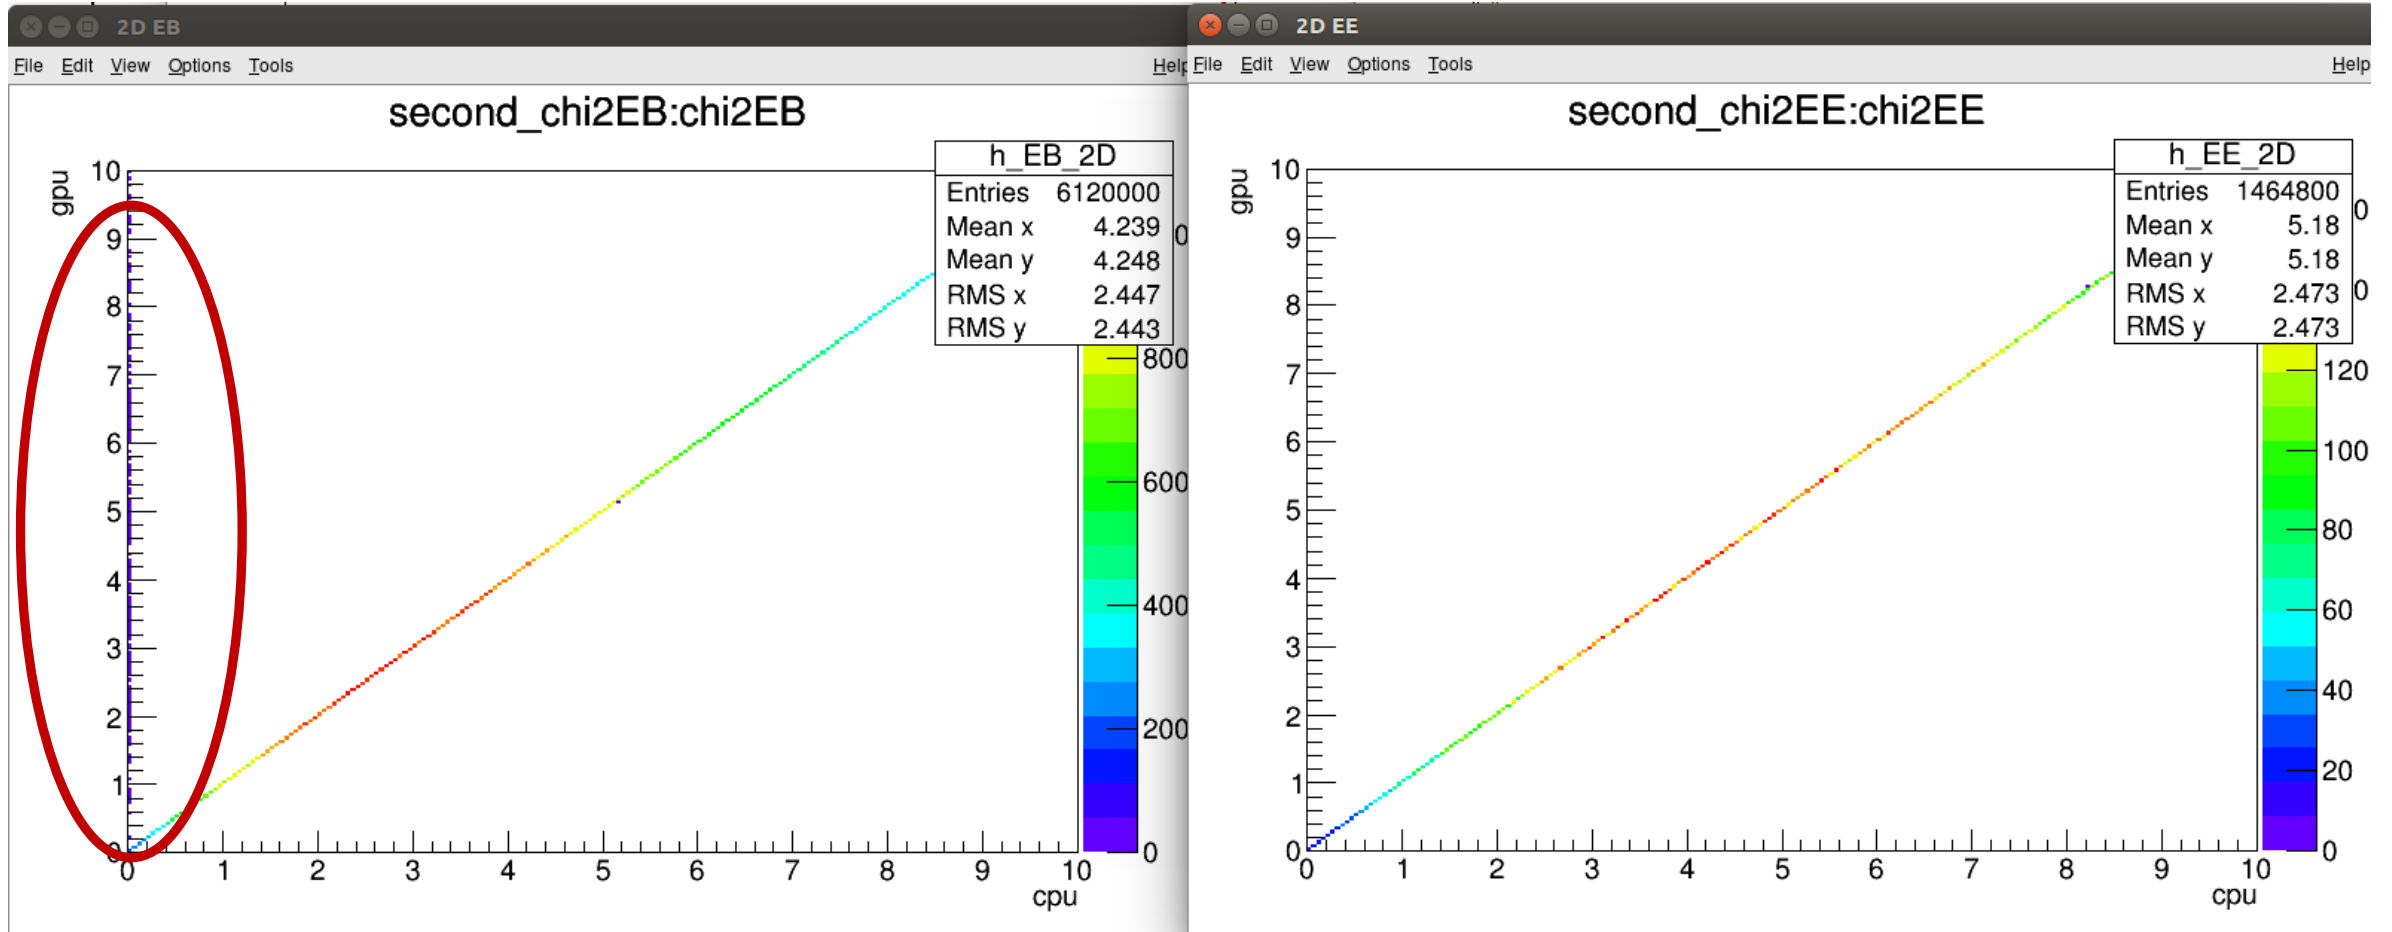
# Ecal Reco Validation (GPU vs CPU)

Reconstructed Energy – good match is observed




# Ecal Reco Validation (GPU vs CPU)

Chi2 – good match



to be resolved

# Ecal Reco Throughput (Work in progress)



Intel Xeon Gold 61XX

- pinned cores to same socket
- 14 cores per socket (no ht)

V100 (1 stream per thread)

P100 (1 stream per thread)

No transfers back

- name
- scan.cpu.workergpu03.02052019.0
  - scan.gpu.workergpu03.p100.02052019.0
  - scan.gpu.workergpu07.v100.07052019.0

There is still room (large)  
For improvement  
(according to Nsight Compute)



# **DEEP** *Projects*



*The DEEP projects have received funding from the European Union's Seventh Framework Programme (FP7) for research, technological development and demonstration and the Horizon2020 (H2020) funding framework under grant agreement no. FP7-ICT-287530 (DEEP), FP7-ICT-610476 (DEEP-ER) and H2020-FETHPC-754304 (DEEP-EST).*