CERN IT & HEP SOFTWARE FOUNDATION

STEFAN ROISER, IT-FTI-PSE

LHC MC WG KICKOFF, 14 NOV 2024

Simulation

Physics Software Engineering

Enrico Bothmann (Jan ’25)
Maksymilian Graczyk
Daniele Massaro
SR
Zenny Wettersten
ORIGIN Fellow (Jan ’25)

CERN IT & HEP SOFTWARE FOUNDATION

STEFAN ROISER, IT-FTI-PSE

LHC MC WG KICKOFF, 14 NOV 2024

CONTENT

- Software engineering and computing infrastructure in collaboration and with a focus on LHC experiments and the MC event generator community
  - Current activities
  - Next steps
  - Discussion on how to broaden the scope further

COMPUTING ASPECTS OF MC EVENT GENERATION AT THE LHC

- Computing needs for upcoming LHC phases may exceed the predicted resource allocations

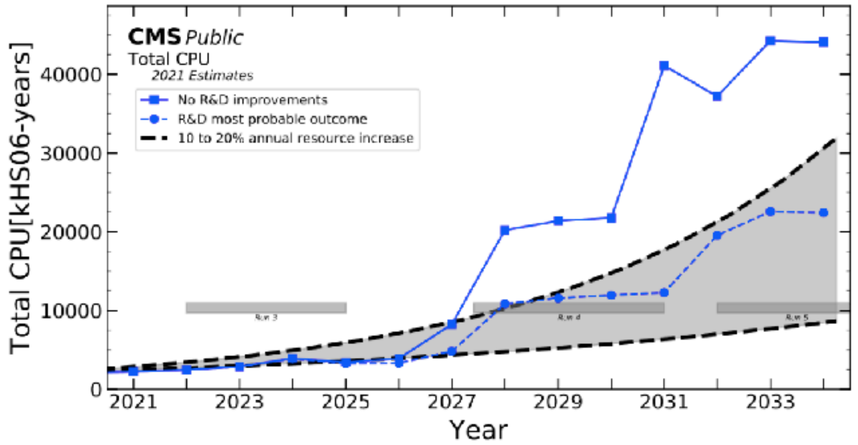

CERN-CMS-NOTE-2022-08

- At the same time the compute time to spend for Monte Carlo event generation will be non-negligible

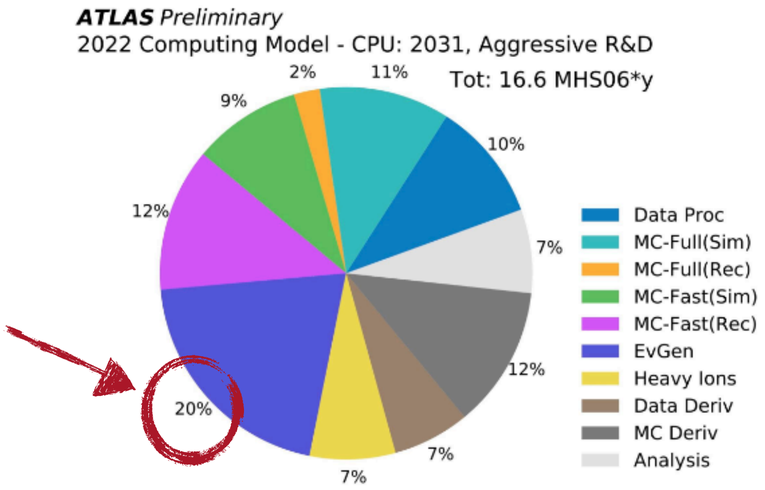

CERN-LHCC-2022-005

STATUS

HEP SOFTWARE FOUNDATION

- Physics generators working group
  - Forum to exchange and gather information: LHC experiments and beyond
    - 27 Nov, seminar on Event generation on GPUs
    - Collecting and preparing input for European strategy for particle physics
- Google summer of code student program
  - Collecting projects and organisation of the program (next round starting in Jan ‘25)
- Training working group on research software engineering topics
  - C++, Python, Julia, ML, ROOT, CMake, git, CI/CD, … see also new training center

ACCESS TO SPECIALISED HARDWARE RESOURCES

https://openlab.cern/heterogeneous-architectures-testbed

CERN openlab

Heterogeneous Architecture Testbed: Hardware

- Euro HPC Supercomputer development access
  - LUMI, Leonardo, Deucalion, Vega, etc

- Networking Switches
  - InfiniBand
  - Intel Omni-Path
- Memory
  - Intel Optane

- CPU
  - Intel Xeon Gen 6, 5, 4 ...
    - Sierra Forest
    - Emerald & Sapphire Rapids
  - IBM Power 8 & 9
  - AMD EPYC*: Genoa, Milan
  - ARM CPUs
    - Fujitsu A64FX**
    - Ampere ALTRA MAX*
    - NVIDIA GRACE
  - RISC-V*

- Accelerators
  - NVIDIA Tesla GPUs
    - L4, T4, A2*
    - A100(X*), V100(S), P100, H100*
  - NVIDIA Bluefield 2 DPU
  - Intel GPUs
    - Max 1100 (Ponte Vecchio)
    - Flex 140 & 170
  - AMD
    - MI100*, MI210x**
    - ALVEO U200 (FPGA)
- AI specialised hardware
  - Cerebras*, Graphcore*

- 100+ users & 290+ accounts
- ~95 systems, mostly bare-metal
- Used by ATLAS, CMS, LHCb, QTI, openlab, CERN EP, IT, ATS
- ~200 tickets handled p/a

*Remote access via E4 COMPUTER ENGINEERING

**Remote access via Simons Foundation

HARDWARE ACCELERATION OF MADGRAPH5_AMC@NLO

- First release for acceleration of leading-order processes on GPUs and vector CPUs available
  - Excellent collaboration with O Mattelaer (Louvain)
  - Already very good speedups of the application via hardware acceleration of matrix elements
  - Available for NVidia and AMD GPUs. SYLC port being worked on
  - Next possible steps for LO: true heterogeneous execution, offloading / acceleration of other parts of the workflow, …

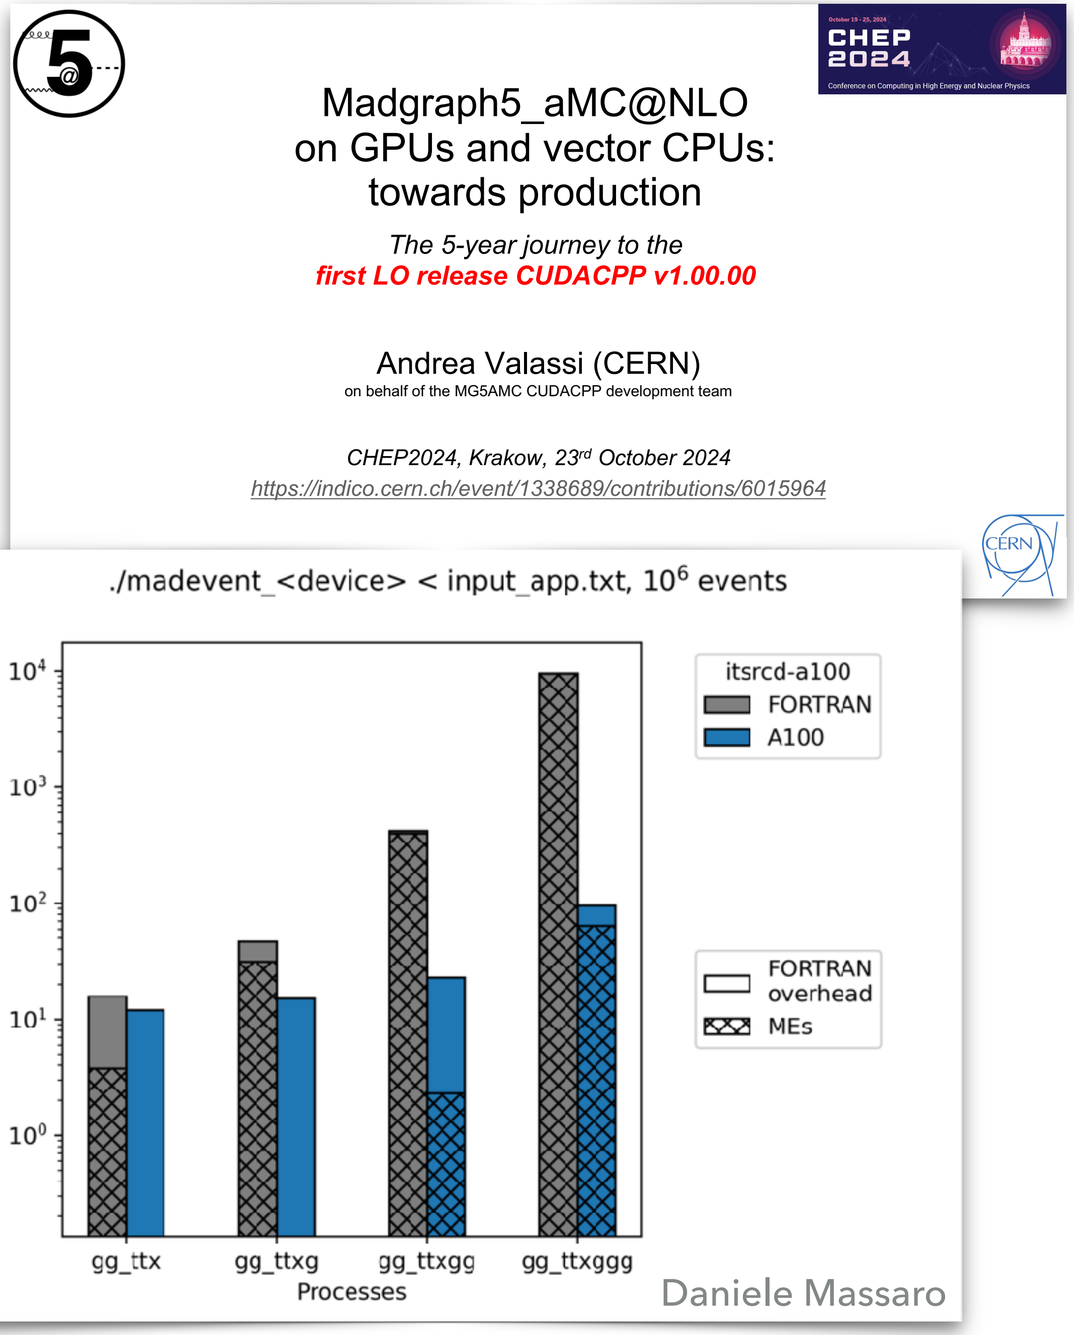

(RE-)USAGES OF MADGRAPH ACCELERATED CODE

- Work on integration with CMS

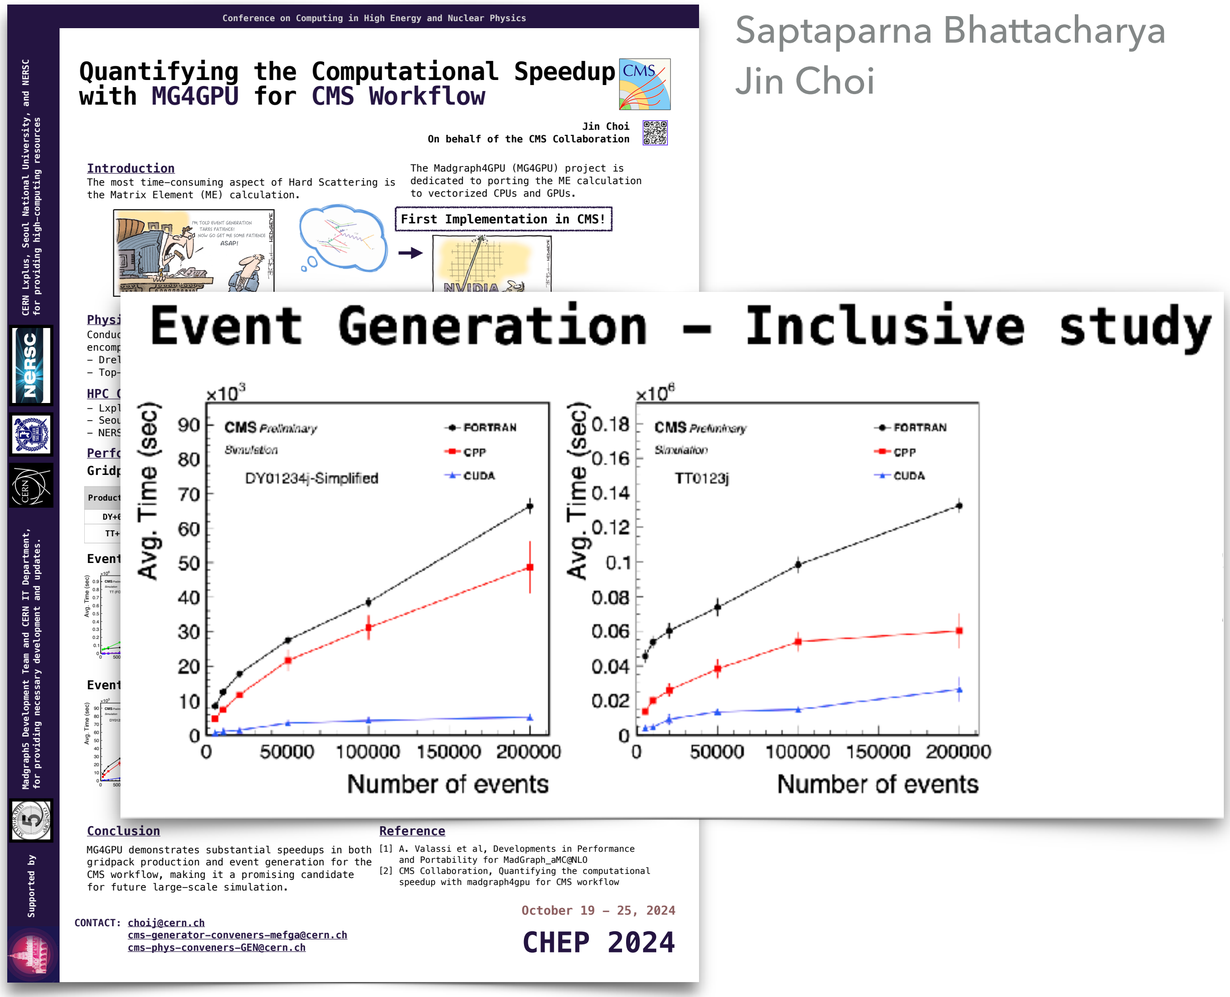

Saptaparna Bhattacharya
Jin Choi

Jin Choi

- More usages of accelerated code
  - Event reweighting

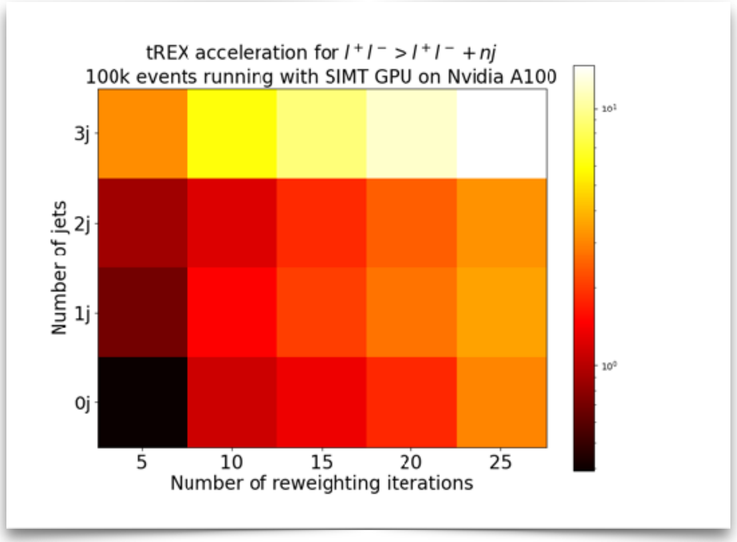

Zenny Wettersten

  - FPGA implementation (U Valencia/CSIC)
  - ME calculations in POWHEG

SHERPA / PEPPER

- Enrico Bothmann joins IT-FTI-PSE in January 2025
- Pepper, GPU code for leading-order parton-level event generation

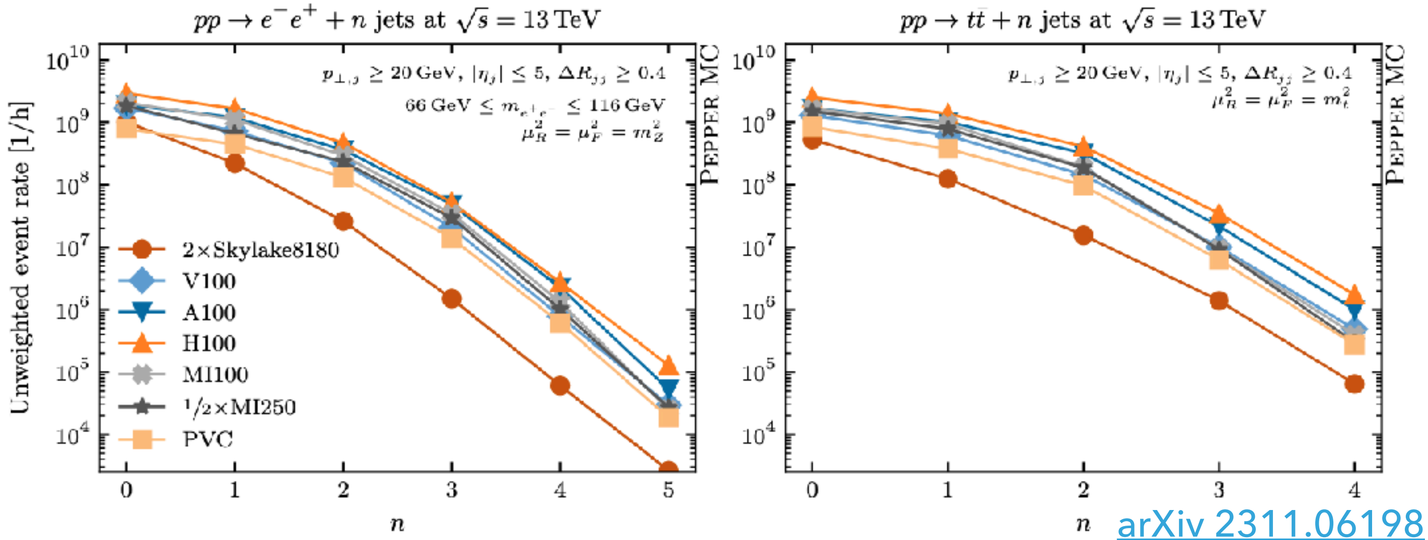

arXiv 2311.06198

  - Integration into experiments workflows ongoing
- Studies performed on large scale (N)LO calculations [arXiv 2309.13154]
  - Combine Sherpa, Pepper & Pythia
  - Addressing I/O bottleneck with HDF5 file format

NEXT STEPS

NEXT-TO-LEADING-ORDER

- Next step for both Madgraph and Pepper
  - Produce prototypes for tree-level calculations
  - Loop calculations will be more difficult

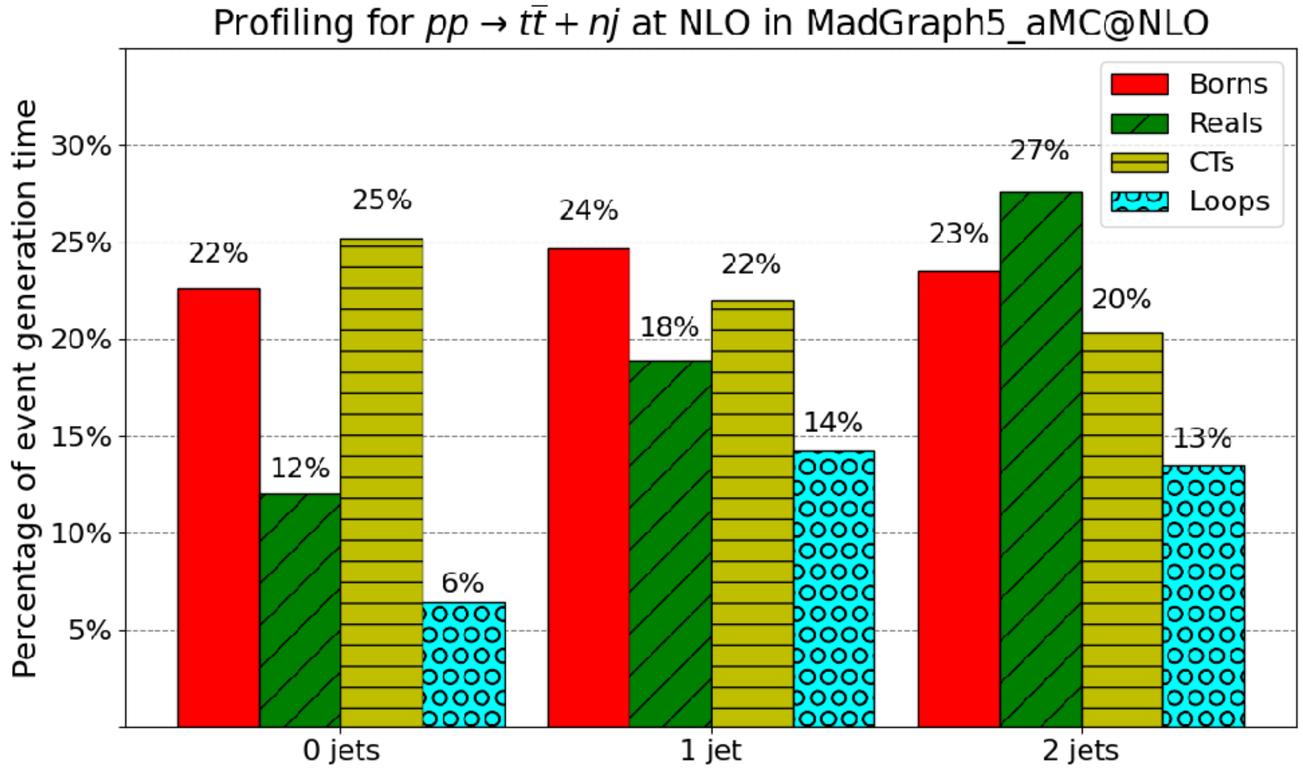

Zenny Wettersten

- How to deal with high precision floating point calculations on future accelerators
  - Double precision FLOPS on GPUs will stay constant at best for next generations
  - How to deal with quad precision, e.g. for loops?

DEVELOP NLO PROTOTYPE FOR HARD-SCATTERING EVENT GENERATION

- Studies and first steps done for Madgraph
- CERN/IT work for Sherpa / Pepper to be defined early next year
- Collaboration with Argonne and Madgraph, Sherpa/Pepper for GPU loop calculations starting

COLLABORATE WITH EXPERIMENTS ON DEPLOYMENT OF ACCELERATED CODES

- Finish up work with CMS for Madgraph
- Work with more experiments and use cases

INVESTIGATE MORE AREAS FOR SOFTWARE ENGINEERING IN MC EVENT GENERATION

- Make more parts of the workflow suitable for hardware accelerated execution
- Interest in going beyond hard scattering, e.g. shower calculations
- Investigation on negative weights

AND MORE . . .

- Double and higher floating point precision calculations on future hardware (PhD)
- Collaborations on student level (AGH/Krakow)
- gitlab.cern for MC event generator projects
- Use of multi-platform event generator code in the context of sustainability studies
- QCD event generation in Julia (HZDR)
- Quantum algorithms (IBM) and QML (arXiv:2409.12236) for event generation

SUMMARY

- Several services, developments and projects in IT and HSF ongoing to support the LHC experiments and the MC event generator community
- Ready to broaden the scope beyond the current activities !!

THANK YOU !!