Community White Paper (CWP) Challenges and Workshop Aims

Elizabeth Sexton-Kennedy, Fermilab
Graeme Stewart, CERN EP-SFT

Physics Event Generator Computing Workshop, CERN, 26 November 2018

HEP Software Foundation

Welcome

A big thank you to everyone for coming!

As stated in the HSF Community White Paper, we know the reward system for developing generator code is not as “suppliers” to the experimental HEP programme.

However, there is a pressing need for exactly that, as it’s a critical part of what we do.

Building a community effort in this area that

- Supports and improves common existing tools like LHAPDF, HepMC, and Rivet
- Encourages use of common tools like phase space integrators
- Facilitates generator physics performance improvements
- Engages engineering help to achieve best technical performance

Software Credit

Some experiments do not recognise contributions to external projects as experiment service work.

CWP suggests reviewing this.

Theorists derive credit from citation of their papers -> requires vigilance

There is recognition that improvement is needed: Credit and recognition for research software: Current state of practice and outlook

CITING GENERATOR TOOLS

On this page we collected citations for the various generators and for some theory paper to be used in CMS papers

Contents:

- CITING GENERATOR TOOLS
  - Recommendations from the Pubcom
  - Info on specific MC campaigns
    - RunII MC (Winter15+Summer15)
  - Theory papers
    - Scale choice and related uncertainty
  - Matrix element generators
    - AlpGen
    - MadGraph5 _aMCatNLO
    - MadGraph5 (prior to the joint MadGraph5_aMCatNLO release)
    - PowHeg
    - Sherpa
  - Parton shower generators
    - Herwig++
    - Pythia6
    - Pythia8
  - Other generators and tools
    - Geant4 (for detector simulation)
    - Fast Sim simulation
    - EvtGen (for B and other hadron decays)
    - PHOTOS (for simulation of QED emission)
    - Tauola (for tau decays)
    - Syscalc
    - Parton distribution functions (for details see PDF group TWIKI)

Recommendations from the Pubcom

Actual references used during final formatting and the preferred macros to use in referring to them are available. For instance, we have \MGvATNLO, which ensures uniformity. The references themselves are in https://svnweb.cern.ch/cern/wsvn/tdr2/utils/branches/dev/general/gen.bib (along with many other common citations). These are taken from the versions published in the journals and sometimes differ from the arXiv or inSPIREHEP versions.

HEP Software Foundation Roadmap for Software and Computing R&D in the 2020s

HSF-CWP-2017-01
December 15, 2017

- HSF established in 2015 to facilitate *coordination* and *common efforts* in software and computing across HEP in general
- Charged by WLCG to address R&D for the next decade
- 70 page document on arXiv (1712.06982)
- **13 topical sections** summarising R&D in a variety of technical areas for HEP Software and Computing
  - Backed by topical papers with more details also (e.g. 50-page detailed review about Detector Simulation)
- **1 section on Training and Careers**
- **310 authors** (signers) from 124 HEP-related institutions
- Feature article in CERN Courier
- More details on the HSF web site

Contents

1 Introduction 2

2 Software and Computing Challenges 5

3 Programme of Work 11
- 3.1 Physics Generators 11
- 3.2 Detector Simulation 15
- 3.3 Software Trigger and Event Reconstruction 23
- 3.4 Data Analysis and Interpretation 27
- 3.5 Machine Learning 31
- 3.6 Data Organisation, Management and Access 36
- 3.7 Facilities and Distributed Computing 41
- 3.8 Data-Flow Processing Framework 44
- 3.9 Conditions Data 47
- 3.10 Visualisation 50
- 3.11 Software Development, Deployment, Validation and Verification 53
- 3.12 Data and Software Preservation 57
- 3.13 Security 60

4 Training and Careers 65
- 4.1 Training Challenges 65
- 4.2 Possible Directions for Training 66
- 4.3 Career Support and Recognition 68

5 Conclusions 68

Appendix A List of Workshops 71

Appendix B Glossary 73

References 79

HEP: Landscape and Frontiers

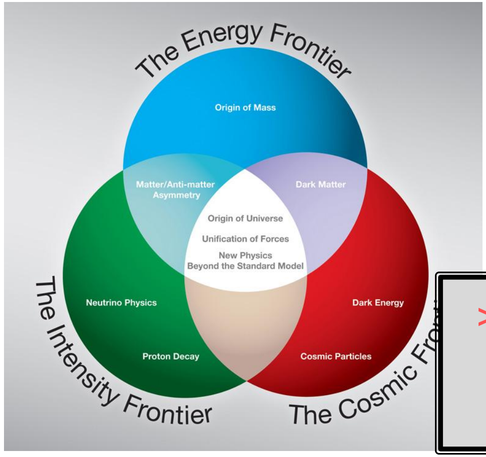

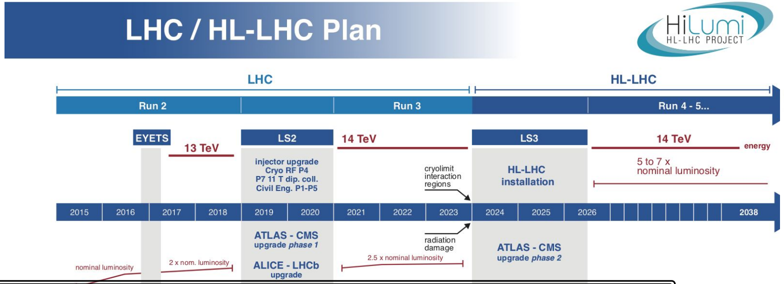

>50M Lines of code that exist in multiple experiments and packages; persistence across generations of experiments

SuperKEKB luminosity projection

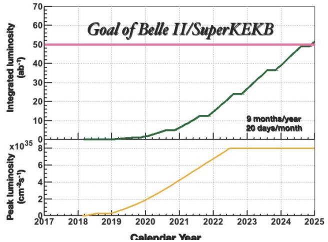

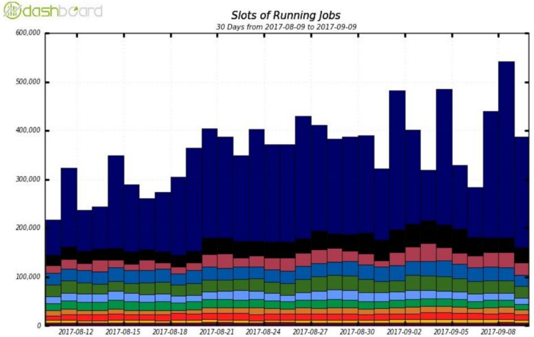

ROOT Data Analysis Framework

HSF HEP Software Foundation

Geant 4

HEP.TrkX

athena

ATLAS Experiment main repository for Athena code

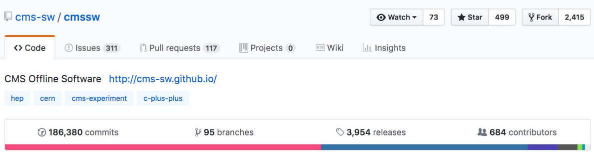

Fermilab Program Planning
20-Feb-17

LONG-RANGE PLAN: DRAFT Version 7a

		FY16	FY17	FY18	FY19	FY20	FY21	FY22	FY23	FY24	FY25	FY26
LBNF/PIP II	LBNF/PIP II							LBNF	LBNF		LBNF / PIP II	
	SANFORD						DUNE	DUNE	DUNE	DUNE	DUNE	DUNE

Summer shutdown | Construction / commissioning | Run

FNAL Intensity Frontier

High Luminosity LHC

- Large rise in rate (~10kHz) and complexity (mu~200) : Run 2 SW & computing will not scale

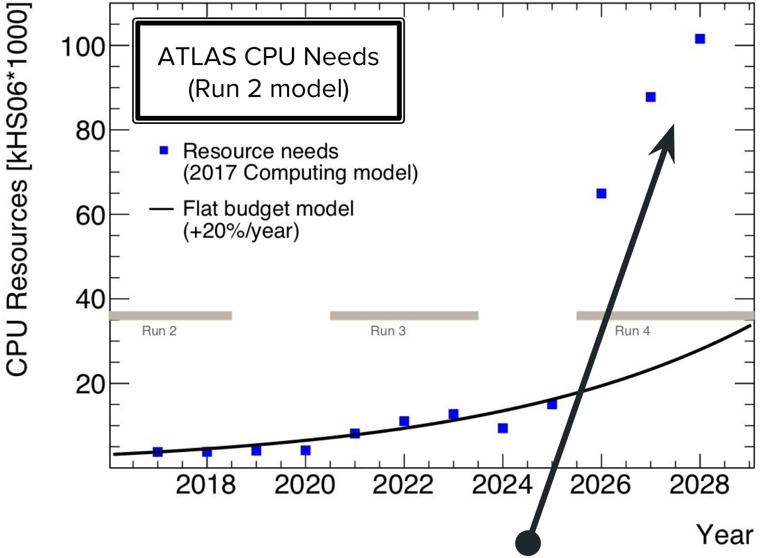

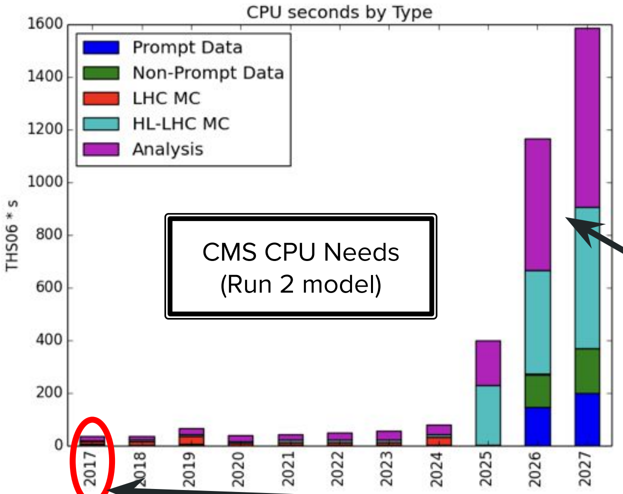

- Resources needed would hugely exceed those from technology evolution alone alone with a flat budget (close to Run 2+3 evolution)

Looking at 2017 in more detail later

Software Challenges for HL-LHC

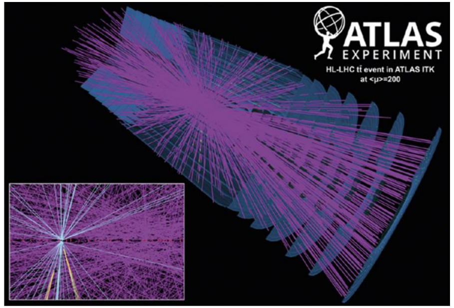

- Pile-up of ~200 ⇒ particularly a challenge for charged particle reconstruction (superlinear scaling, ~x30-50)
- With a flat budget, improvements from hardware of ~x6 (Moore's Law) are the real maximum we can expect
- Increased amount of data requires us to revise/evolve our computing and data management approaches
  - We must be able to feed our applications with data efficiently at scale (end-to-end computing)
  - For analysis sheer volume of event data is a major factor - I/O bound workload
- HEP software typically executes 1 instruction at a time (per thread)
  - Major re-engineering required to benefit from modern CPUs (can do 8 in theory, more like 2-4 for 'real' code)
  - Accelerators like GPUs are even more challenging
- HL-LHC salvation will come from software improvements, not from hardware

HL-LHC Computational Intensity

The “physics of interest” is at smaller cross section AND higher final state particle multiplicity.

Note the pp inelastic cross section doesn’t even fit on this log scale plot …another source of computational complexity

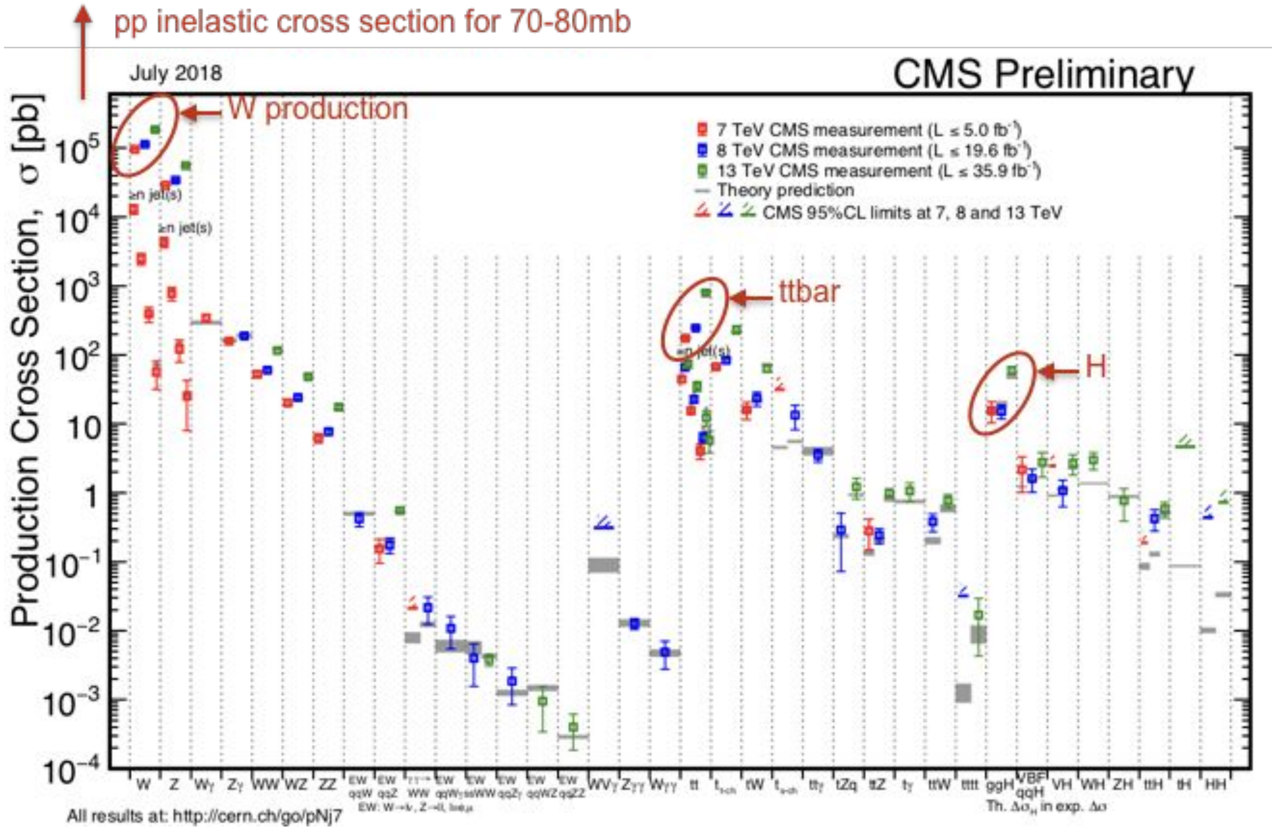

Generators CAN be Computationally Intensive

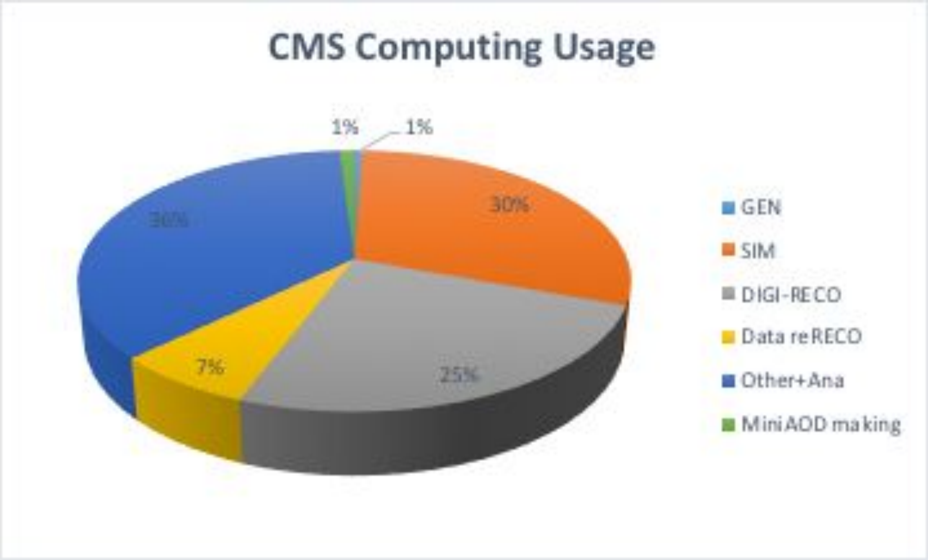

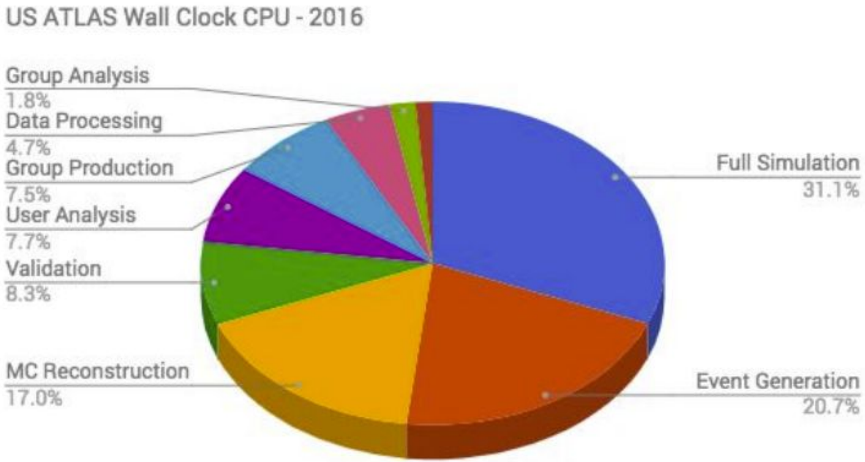

CMS usage from 2017, ATLAS went down to 14.3% in 2017

- These values vary from year to year as analysis needs vary
- CMS uses more LO samples in this year and grid-pack configurations

Challenges at Next-to-Leading Order

- NLO Generators (aMC@NLO/Sherpa) are also widely used for LHC studies.
- Computational intensity grows roughly factorially with number of particles and virtual loops. O(100k) diagrams at LO for W+5jets, O(100k) at NLO for W+3jets
- All possible diagrams must be represented in memory during event generation
- Virtual loops also drive compute intensity and in rare cases require quad precision due to large cancellations between individual diagrams.
- Currently pair hard scatter with LO showering, but NLO showering is being developed and increases computational intensity further.

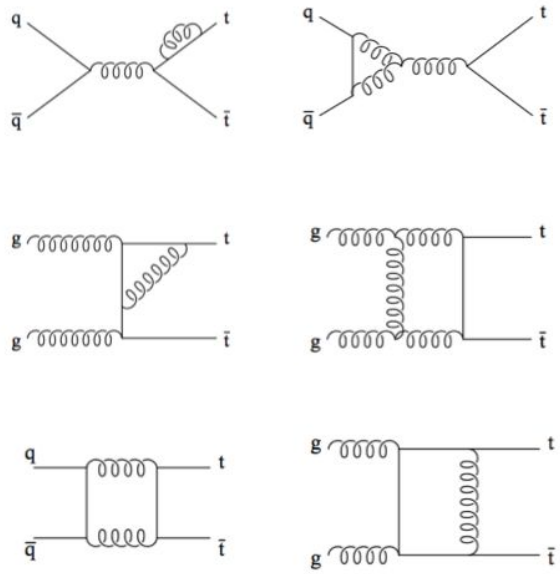

Going Boldly Beyond NLO

- ATLAS W+jets production, leading jet pT shows we are already in the NNLO regime,
- Theoretical methods for calculating high multiplicity processes do not yet exist at NNLO with the current limit being 2 jets
- Theorist community and MC development community continues to innovate and push the orders of perturbative QCD
- HL-LHC we will drive the need for more MC events at high order to support high precision studies across high-dimensional phase spaces.
- HPCs are becoming a focus of the community, as this does become a genuine computing problem: high-dimensional phase-space require more MC integrand evaluations.

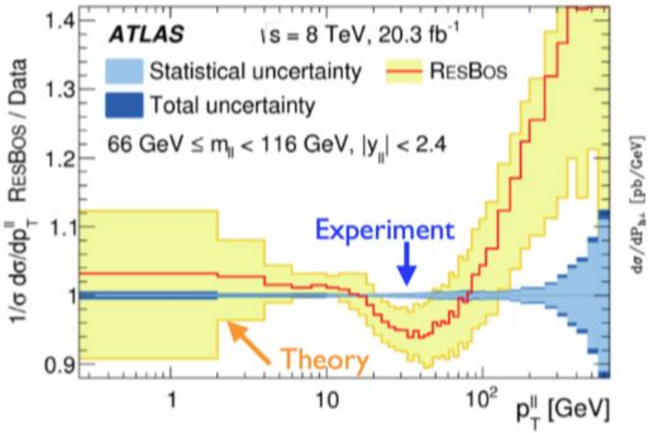

Theory errors more than an order of magnitude larger than experimental ones

https://indico.cern.ch/event/557731/contributions/2268995/attachments/1342762/2022840/Boughezal-HPC2016-Sep22.pdf

Recap: Run 1 -> Run 2 -> Run 4?

1. In Run 1 ATLAS/CMS relied on LO generators which took O(100ms) to execute
2. In Run 2 we moved to NLO generators with O(10-100s) execution times
3. The total grid times of each experiment depend on the mixture of 1. + 2. That their analyzers request.
   a. This creates a large dynamic range in resources needed for samples
   b. Introduces uncertainty in computing needs
   c. Creates lots of validation headaches
4. ATLAS/CMS create O(10B events) / year. Rule of thumb 1-2x collected events
   a. Creates a need for robustness and performance
   b. Can't expect to generate these samples 100 events at a time
5. Run 4 will:
   a. Collect 10x Run 2 number of events
   b. More precision -> higher order
   c. More particles, bigger phase spaces
6. All within a FLAT computing budget

What does Modernization Mean?

The computing available in 2026 will be heterogeneous and highly concurrent.

Our physics reach will be limited by computing budgets so we must optimize for this computing landscape including taking full advantage of:

1. Many Core threading models
2. Single Instruction Multiple Data (SIMD) vectorization
3. Non-uniform Memory Access (NUMA) hierarchies
4. Offloading to accelerators like
    a. Graphic Processing Unit (GPU),
    b. Field Programmable Gate Array (FPGA),
    c. and Tensor Processing Unit (TPU)

A move towards an Open Source development model is a must for the theory community. This is standard practice for most (all?) experiments.

Addressing the Challenges - Day 1 Morning

Given the mixed audience we thought we'd have an experimental introduction and a generator intro. describing the different moving parts of generator codes.

These talks will focus on experiment requirements. For example, what is needed in terms of accuracy & statistics so as not to be limited by MC?

Monday, 26 November

09:00	→ 09:05	**Welcome and Practicalities** Speakers: Andrea Valassi (CERN), Graeme Stewart (CERN)	5m	4-3-006 - TH Conference Ro...
09:05	→ 09:30	**CWP Challenges and Workshop Aims** We review the challenges for the HEP field with an emphasis on Physics Generators as an introduction to the workshop's goals Speakers: Elizabeth Sexton-Kennedy (Fermi National Accelerator Lab. (US)), Graeme Stewart (CERN)	25m	4-3-006 - TH Conference Ro...
09:30	→ 10:00	**Introduction to Generator Codes** Speaker: Simon Platzer (University of Vienna (AT))	30m	4-3-006 - TH Conference Ro...
10:00	→ 10:20	**ATLAS Needs and Concerns** Speaker: Josh McFayden (CERN)	20m	4-3-006 - TH Conference Ro...
10:20	→ 10:40	**CMS Needs and Concerns** Speaker: Efe Yazgan (Chinese Academy of Sciences (CN))	20m	4-3-006 - TH Conference Ro...
10:40	→ 11:10	Coffee	30m	4-3-006 - TH Conference Ro...
11:10	→ 11:30	**LHCb Needs and Concerns** Speaker: Philip Ilten (University of Birmingham (GB))	20m	4-3-006 - TH Conference Ro...
11:30	→ 11:50	**Practical Computing Considerations** Speaker: David Lange (Princeton University (US))	20m	4-3-006 - TH Conference Ro...
11:50	→ 12:15	**Discussion**	25m	4-3-006 - TH Conference Ro...

Addressing the Challenges - Day 1 Afternoon

Here we hear in more detail about the technical requirements on modern codes, suitable for use on HPCs.

Status and plans of the widely used generators. We hope to have a focus on algorithms and performance. Is it possible to define an agreed upon benchmark physics process?

12:15 → 13:15	Lunch	1h	4-3-006 - TH Conference Ro...
13:15 → 13:45	Performance Optimisation Introduction Speaker: Servesh Muralidharan (CERN)	30m	4-3-006 - TH Conference Ro...
13:45 → 14:15	Optimising Memory Use Speaker: Sebastien Ponce (CERN) MemoryUsageO...	30m	4-3-006 - TH Conference Ro...
14:15 → 15:00	Performance Discussion	45m	4-3-006 - TH Conference Ro...
15:00 → 15:30	Coffee	30m	4-3-006 - TH Conference Ro...
15:30 → 15:53	Sherpa Status and Plans Speaker: Marek Schoenherr (CERN)	23m	4-3-006 - TH Conference Ro...
15:53 → 16:16	Madgraph Status and Plans Speaker: Olivier Mattelaer (UCLouvain)	23m	4-3-006 - TH Conference Ro...
16:16 → 16:38	PYTHIA 8 Status and Plans Speaker: Philip Ilten (University of Birmingham (GB))	22m	4-3-006 - TH Conference Ro...
16:38 → 17:00	Herwig Status and Plans Speaker: Peter Richardson (CERN & IPPP, Durham University)	22m	4-3-006 - TH Conference Ro...
17:00 → 17:20	POWHEG Status and Plans Speaker: Emanuele Re (CERN)	20m	4-3-006 - TH Conference Ro...
17:20 → 18:00	Discussion	40m	4-3-006 - TH Conference Ro...

Addressing the Challenges - Day 2 Morning

Reports of ongoing work to make generator codes more suitable for use on HPC systems.
Can these ideas be made suitable for many generators?

Event reweighting?
Sharing between ATLAS & CMS ?

Status/plans for generators used as “decayers”

Tuesday, 27 November

09:00	09:25	HDF5 usage for HEP event generation	Speaker: Holger Schulz (Fermilab)	25m	4-3-006 - TH Conference Ro...
09:25	09:50	A novel workflow of generator tunings in HPC for LHC new physics searches	Speaker: Xiangyang Ju (Lawrence Berkeley National Lab. (US))	25m	4-3-006 - TH Conference Ro...
09:50	10:15	Adaptive multi-channel integration with MPI	Speakers: Simon Brass (University of Siegen), Simon Braß (Universität Siegen)	25m	4-3-006 - TH Conference Ro...
10:15	10:30	Discussion		15m	4-3-006 - TH Conference Ro...
10:30	11:00	Coffee		30m	
11:00	12:00	Optimising Generator Useage	Speakers: Andy Buckley (University of Glasgow (GB)), Marek Schoenherr (CERN), Stefan Hoeche (SLAC)	1h	4-3-006 - TH Conference Ro...
12:00	12:15	Photos and Tauola Status and Plans	Speaker: Zbigniew Andrzej Was (Polish Academy of Sciences (PL))	15m	4-3-006 - TH Conference Ro...
12:15	12:30	EvtGen Status and Plans	Speaker: Michal Kreps (University of Warwick (GB))	15m	4-3-006 - TH Conference Ro...

Addressing the Challenges - Day 1 Morning

Beyond current practice

Follow up planning

Hackathon, sign up here:
https://docs.google.com/spreadsheets/d/1uReZEVU-wS0IiNiqaRjci7fPrJpBPtzrfaa9CKfTbTU/edit#gid=0

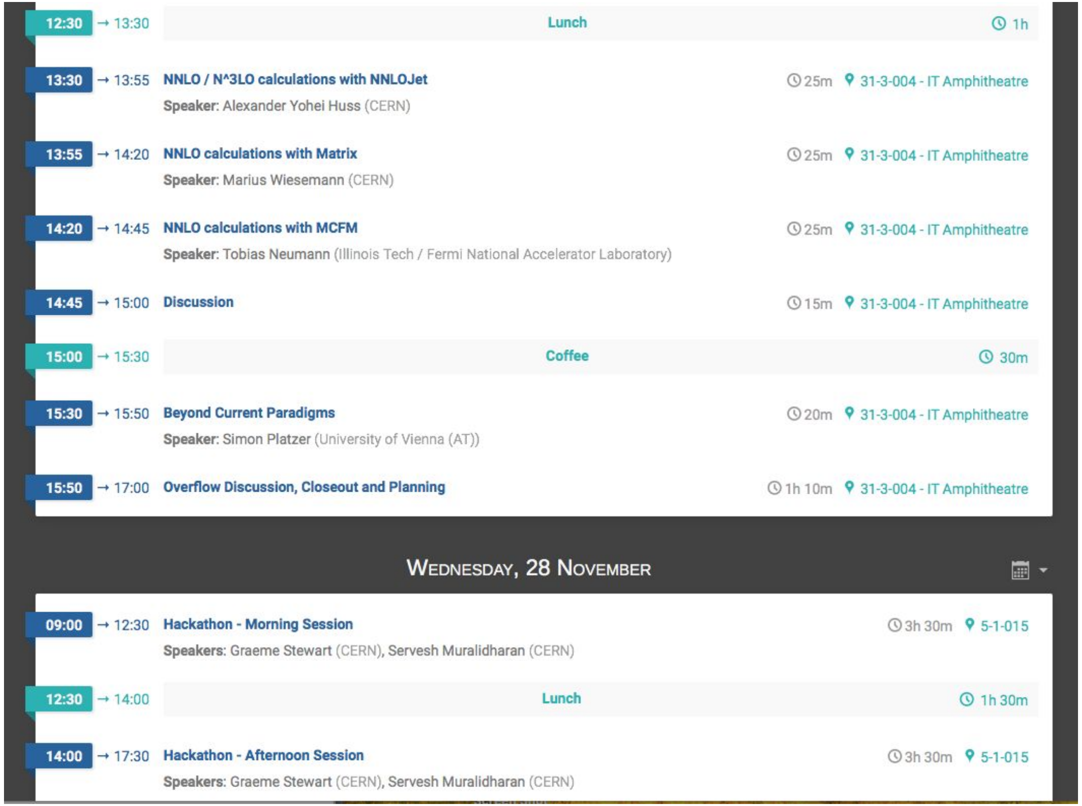

12:30 → 13:30 Lunch 1h

13:30 → 13:55 NNLO / N^3LO calculations with NNLOJet 25m 31-3-004 - IT Amphitheatre
Speaker: Alexander Yohei Huss (CERN)

13:55 → 14:20 NNLO calculations with Matrix 25m 31-3-004 - IT Amphitheatre
Speaker: Marius Wiesemann (CERN)

14:20 → 14:45 NNLO calculations with MCFM 25m 31-3-004 - IT Amphitheatre
Speaker: Tobias Neumann (Illinois Tech / Fermi National Accelerator Laboratory)

14:45 → 15:00 Discussion 15m 31-3-004 - IT Amphitheatre

15:00 → 15:30 Coffee 30m

15:30 → 15:50 Beyond Current Paradigms 20m 31-3-004 - IT Amphitheatre
Speaker: Simon Platzer (University of Vienna (AT))

15:50 → 17:00 Overflow Discussion, Closeout and Planning 1h 10m 31-3-004 - IT Amphitheatre

Wednesday, 28 November

09:00 → 12:30 Hackathon - Morning Session 3h 30m 5-1-015
Speakers: Graeme Stewart (CERN), Servesh Muralidharan (CERN)

12:30 → 14:00 Lunch 1h 30m

14:00 → 17:30 Hackathon - Afternoon Session 3h 30m 5-1-015
Speakers: Graeme Stewart (CERN), Servesh Muralidharan (CERN)

Summary

HEP Software Foundation

- Major challenges for software and computing come in the future
- HL-LHC physics analysis results require software tooling that will deal with a huge increase in events, as driven by physics and the resulting huge increases in generated events.
- Goals of this workshop:
    - Identify the most crucial areas for technical improvements to the generators used by the experiments.
    - Define a programme of work that can be used to attract investment in these technical areas, aiming to have software engineers who can work together with the generator authors.
    - Identify ways of making new theoretical advances easier to implement in a computationally efficient way.

Acknowledgements

Many thanks to:

Graeme Stewart, Taylor Childers

And to the whole community for their excellent work in HEP Software and Computing