LHAPDF Status report

Max Knobbe

MC Support Tools Workshop @ IPPP Durham

17.03.2026

LHAPDF for HPC applications

- Multiple groups working on workflows beyond the traditional CPU + RAM setup
  - LHAPDF usually not time critical, but significantly simplifies workflows if available for e.g. GPUs
- Developed sometime during pandemic, employment in the last years

Testing on multiple architectures

- Want to stay flexible and not bank on solely one framework
  - Keep Cuda and Kokkos
- Tested on multiple architectures at Argonne National Lab
  - NVIDIA, AMD, INTEL,…
- Looking good, testing WIP

```
NVIDIA A100-PCIE-40GB, 40960 MiB, 590.48.01
--------------------------------------------------------------------------
Starting benchmark at Mon Mar 16 03:57:32 PM UTC 2026
LHAPDF 6.5.3 loading /home/ac.mknobbe/lhapdf-sets/CT10nlo/CT10nlo_0000.da
CT10nlo PDF set, member #0, version 4; LHAPDF ID = 11000
LHAPDF 6.5.3 loading /home/ac.mknobbe/lhapdf-sets/CT10nlo/CT10nlo_0000.da
CT10nlo PDF set, member #0, version 4; LHAPDF ID = 11000
LHAPDF 6.5.3 loading /home/ac.mknobbe/lhapdf-sets/CT10nlo/CT10nlo_0000.da
CT10nlo PDF set, member #0, version 4; LHAPDF ID = 11000

Performance benchmark: CT10nlo member 0
Repetitions per batch size: 1
--------------------------------------------------------------------------
    N points      GPU (ms)      CPU (ms)       Speedup    GPU/point (us)
--------------------------------------------------------------------------
          64         0.043         0.022          0.5x             0.674
         128         0.043         0.032          0.8x             0.334
         256         0.044         0.056          1.3x             0.171
         512         0.044         0.110          2.5x             0.086
        1024         0.044         0.218          5.0x             0.043
        2048         0.045         0.432          9.7x             0.022
        4096         0.045         0.895         20.1x             0.011
        8192         0.045         1.776         39.7x             0.005
       16384         0.052         3.429         65.4x             0.003
       32768         0.064         6.843        107.8x             0.002
       65536         0.087        13.744        158.7x             0.001
      131072         0.157        27.566        175.1x             0.001
      262144         0.269        56.199        208.9x             0.001
--------------------------------------------------------------------------
```

Deployment:

- Used in 2 larger campaigns:
  - For DY project in the UK: roughly 50k GPUhrs — no issues
  - Scaling tested as part of Pepper on large US machines:
    - Looking good as well

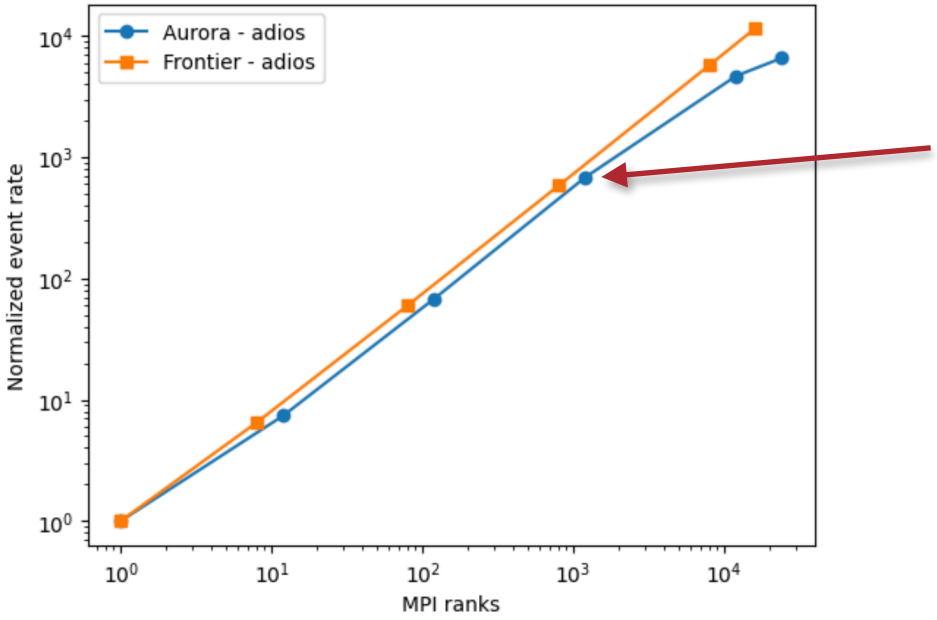

WLCG

Fermilab

LHAPDF PDF submission

- Essentially a three step procedure:
    1. Receive PDF, verify metadata and grids
    2. Assign unique id
    3. Sync with central storage at CERN and update website

1. In the past, via email; easy to miss, forget
    1. AB moved the system to gitlab for simpler tracking 🍀
    2. CG created website for upload of PDFs 🍀
    3. Validation now automatic
2. Manual assignment, group by collab, type of PDF
3. Also manually done at CERN

Room for improvement?

Conclusions

- LHAPDF not in a horrible state
  - Some hiccups with automatic submission procedures but process stabilizing and hopefully much more convenient

- Ready for different computing architectures and future machines: used in production on Tursa@Edinburgh and Aurora@ANL

- Multi PDF access currently being integrated and tested, performance looking promising

⇒ LHAPDF-7 soon?!

FERMILAB-CONF-25-0448-T

Fermilab U.S. DEPARTMENT OF ENERGY Office of Science

NP calibration with History matching

Max Knobbe

MC Support Tools Workshop @ IPPP Durham

17.03.2026

Experimental Relevance

Slide from Jennifer Roloff @ PSR 2025

jet calibrations

- Obvious trend in the jet response from the hadronization model
  - *Calorimeter response depends on the type of hadron, not just the energy and rapidity*
  - Using **retuned Sherpa** with LEP data on baryon and kaon fractions has a significant effect on the p_T response!

ATL-PHYS-PUB-2022-021

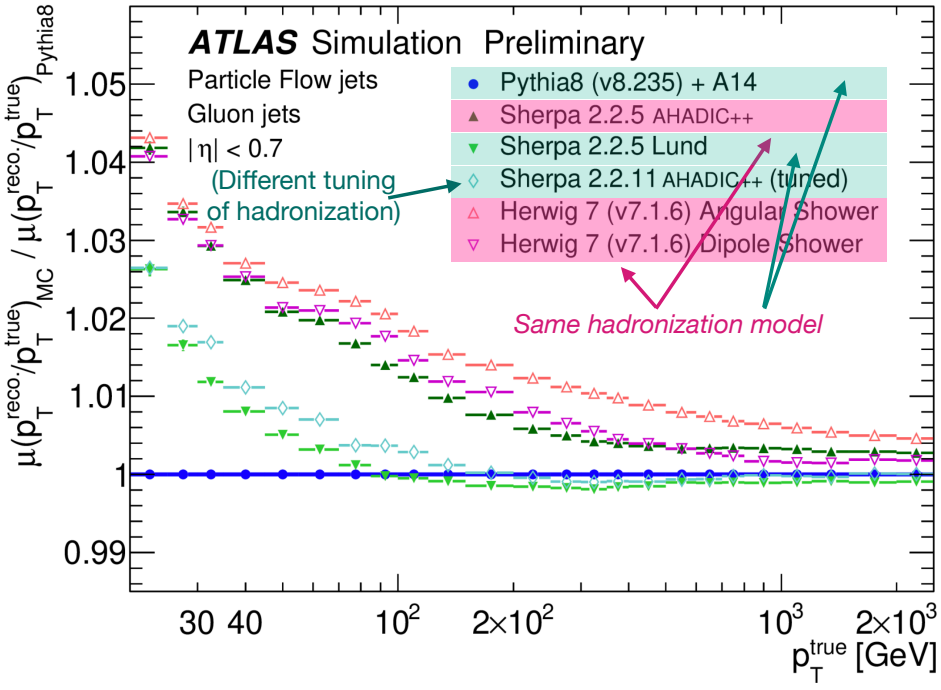

Fermilab

NP Variations

- Comparing different NP models and tunes
  - Sherpa + Cluster/String creates relatively small variations (5-10%)
  - Largest variations between different Pythia tunes, but also not huge

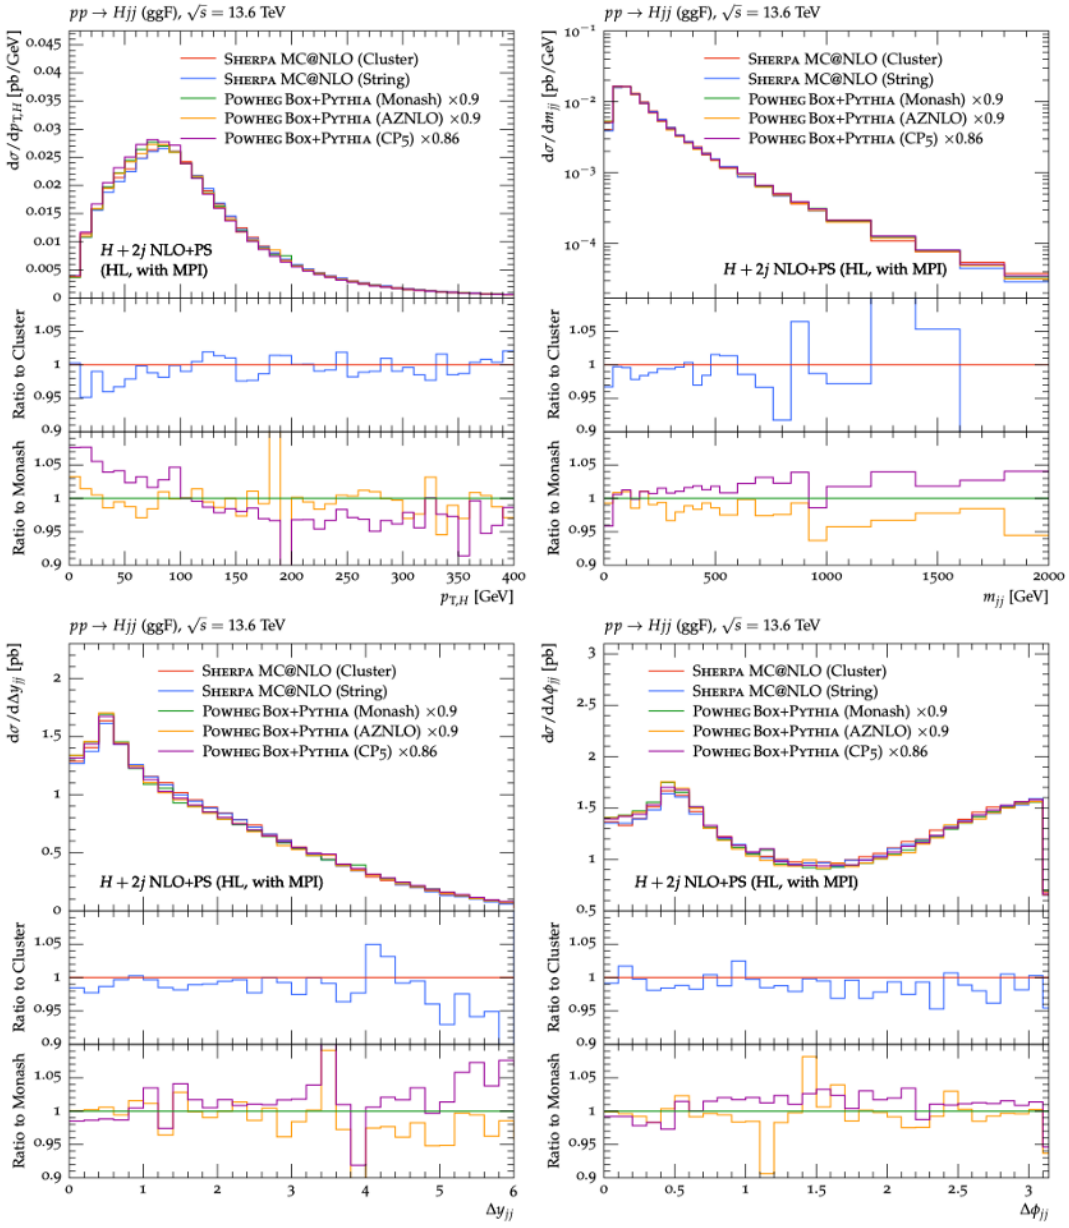

Fermilab

Tuning: Setting the scene

Established workflows:

- Use combination of Professor/Appretice + Rivet
  - Fast surrogate for minimisation, derived from grids of runs

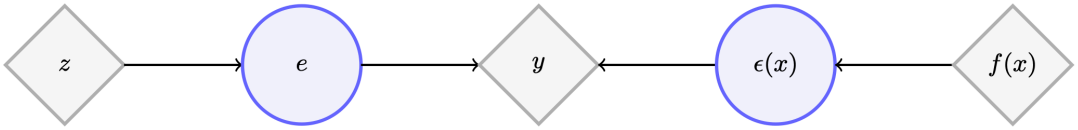

- In Sherpa: Model processes at NLO+PS where possible/ applicable
- Initial scan of wide parameter intervals using 1000s or runs
- Iteratively narrowing ranges and re-preducing grids and surrogates
- Crucially: dedicated campaigns for separate physics models

New dev: calibration w. History Matching

[Iskauskas, MK, Krauss, Schumann, 2602.2232]

- HM identifies non-implausible parameter regions of parameter space
- Observation z of physics process y, simulator prediction given by $f(x)$, e, $\epsilon(x)$ observational errors, model disc.

  - Measure to exclude *implausible* regions

$$I^2(x) = \frac{E(f(x) - z)^2}{Var(f(x) - z)} \approx \frac{(f(x) - z)^2}{Var(e) + Var(\epsilon(x))}$$

  - Consecutive compression of valid volume (HM waves)
- Approximate $f(x)$ with Bayes Linear Emulator

New dev: calibration w. History Matching

[Iskauskas, MK, Krauss, Schumann, 2602.2232]

- Applied to Sherpa hadronisation tuning for $e^-e^+ \to$ Hadrons for internal cluster model and Lund string via Pythia8
- Global LEP callibration of 20/23 parameters in 3/5 waves
- Each wave starting with 1000 points, later reduced to 800

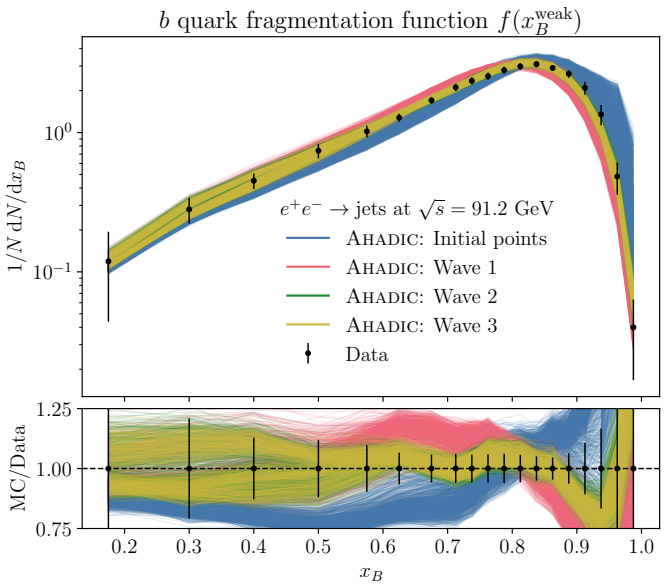

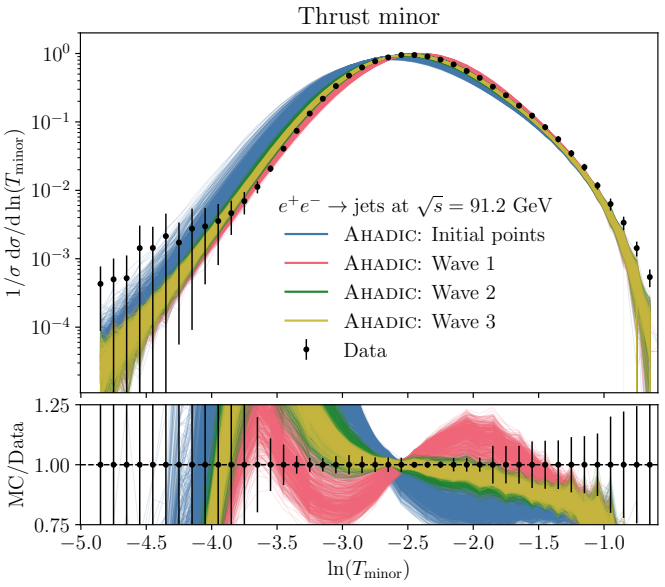

Fermilab

New dev: calibration w. History Matching

[Iskauskas, MK, Krauss, Schumann, 2602.2232]

- Read final wave as estimate on NP uncertainties

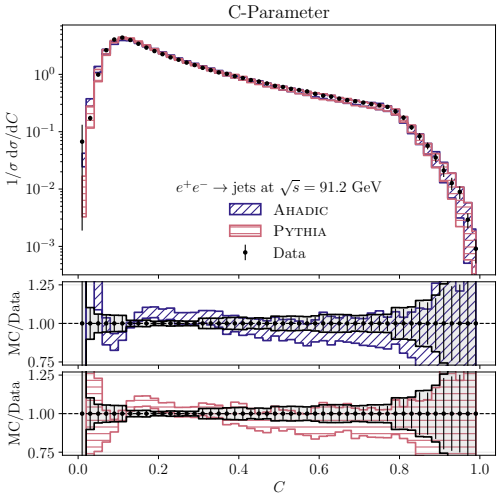

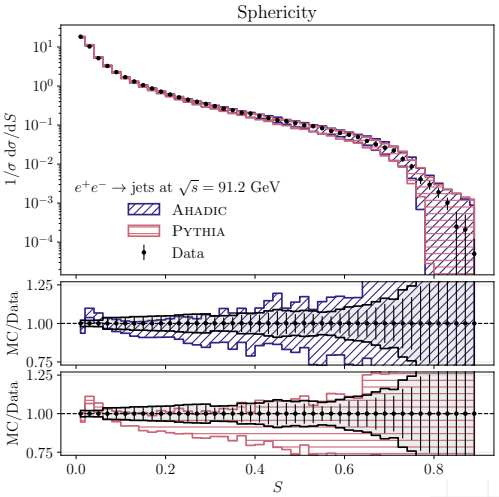

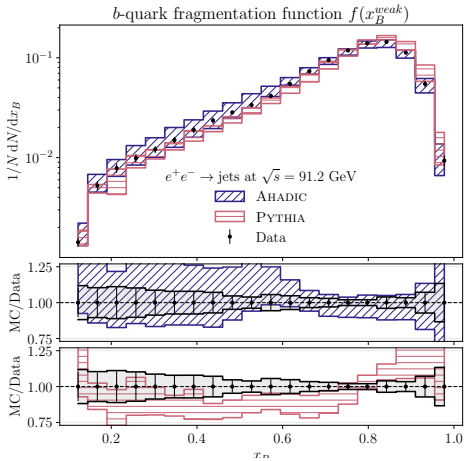

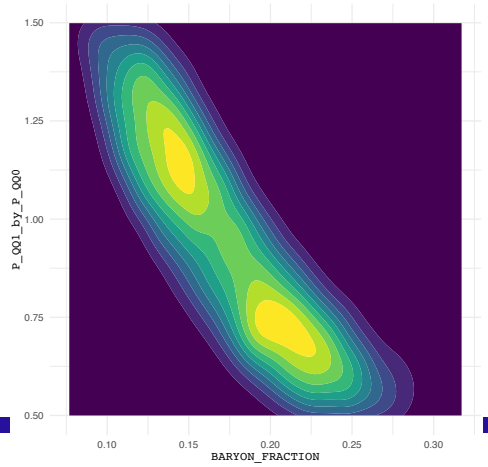

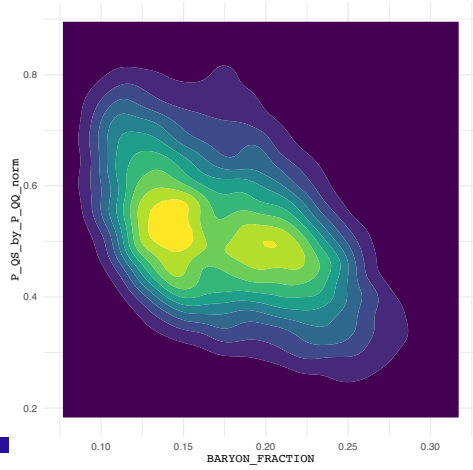

Fermilab

But how to eff. deal with variations?

[Buckley, MK, Krauss, Schumann, TBP , and MK Phd Thesis]

- On the fly reweighing of Hadronisation, MPI and CR params available soon
- Results in alternative event weights, as for perturbative variations (PDFs, scales, ..) [For Sherpa: Bothmann et al. EPJC 76 (2016) 11]
- Allows to efficiently pass through param variations and crucially does not require additional detector simulation

Reweighting: the Hadronisation case

[Buckley, MK, Krauss, Schumann, TBP , and MK Phd Thesis]

- Able to reweight all hadronisation parameters: flavours and kinematics
-

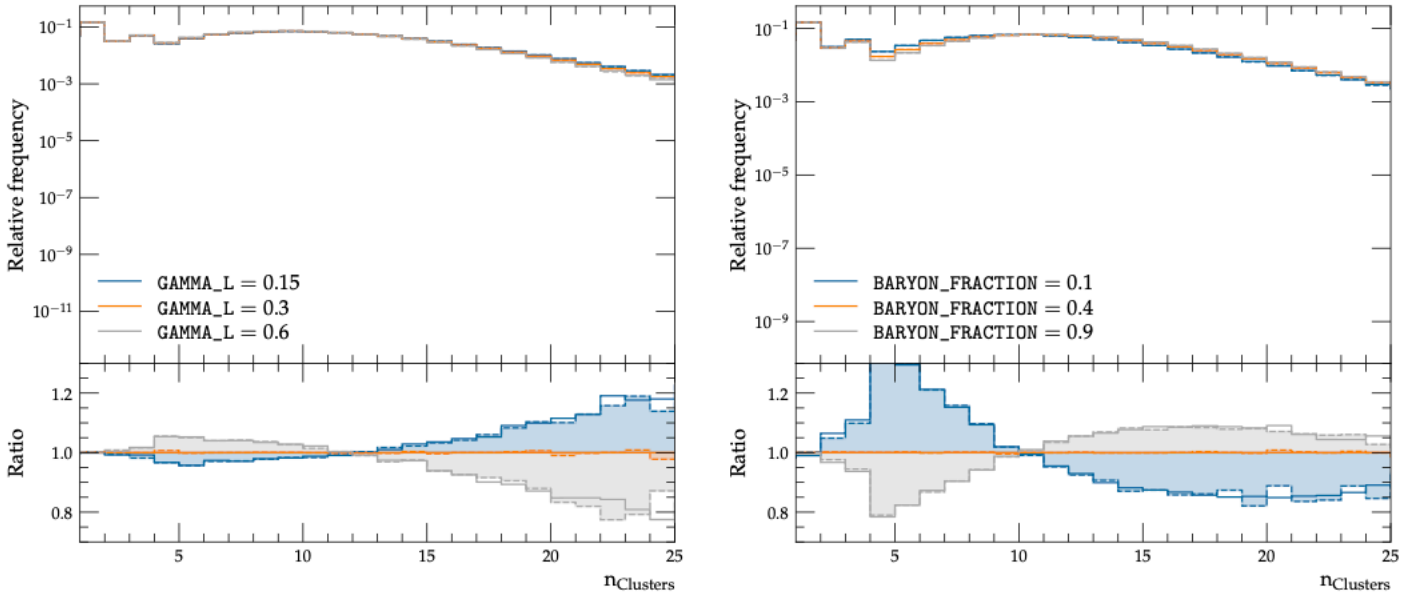

Fermilab

Reweighting: the CR/MPI case

[Pabst, MK, Krauss, Schumann, TBP]

- Sherpa comes with handful of parameters for MPI and CR
  - Performed extreme reweighing test: complete allowed range
- Reweighting works, hence reweighing of any parameter set should also work

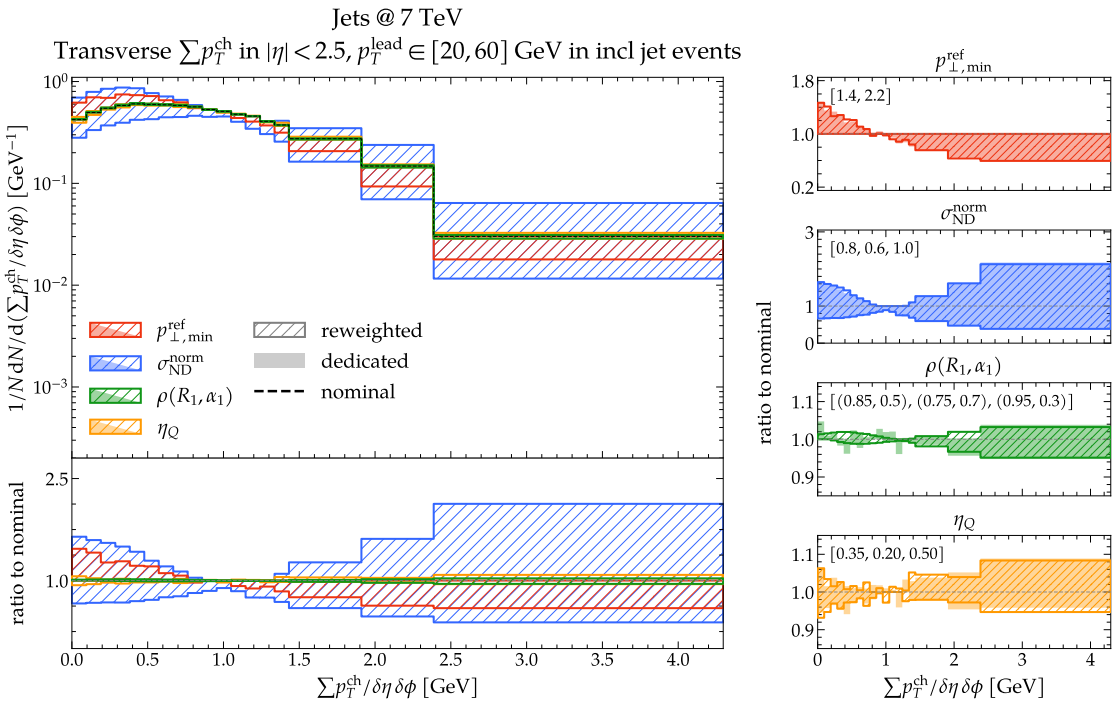

Fermilab

Reweighting: the CR/MPI case

[Pabst, MK, Krauss, Schumann, TBP]

- Similarly for CR reweighing

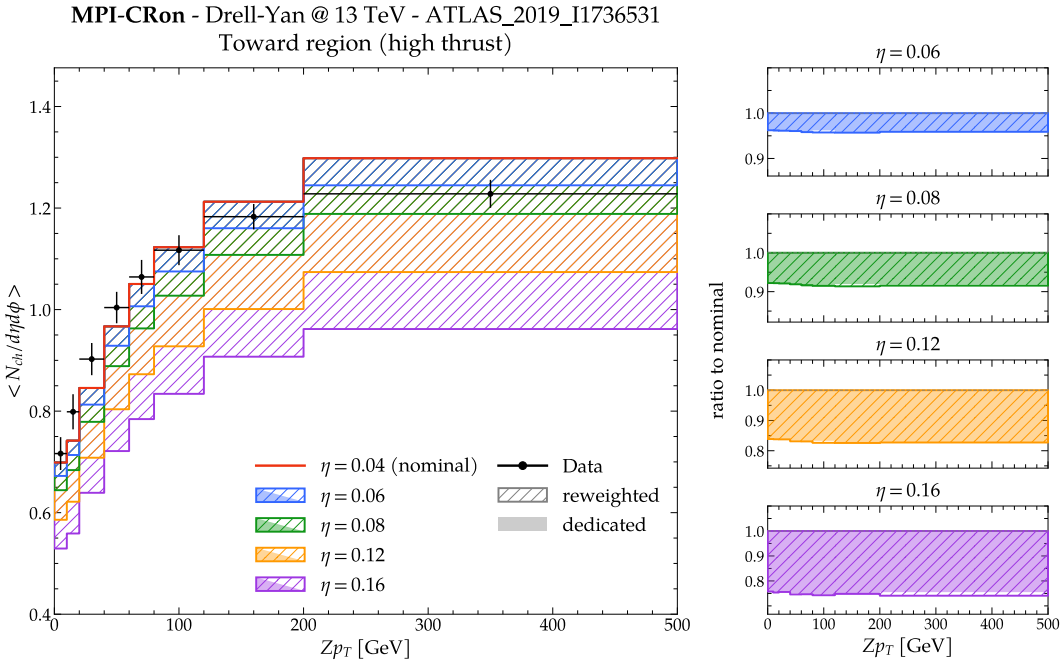

Fermilab

Conclusions

- Synergetic new tuning methodology
  - Requires reweighing techniques to make use of the full potential

- In Sherpa: soon reweighing of Hadronisation, MPI and CR NP parameters

- Potential for number of follow up studies; accelerated tuning workflows, joint tuning of different models, … ?