# Next-To-Next-To-Leading Order Calculations on HPC

**Radja Boughezal**

Argonne NATIONAL LABORATORY

Beyond Leading-order Calculations on HPC

September 22, 2016

# How we make predictions

- We expand our cross section in powers of the strong coupling constant **$\alpha_s$** (a perturbative expansion), and calculate each term in that expansion:

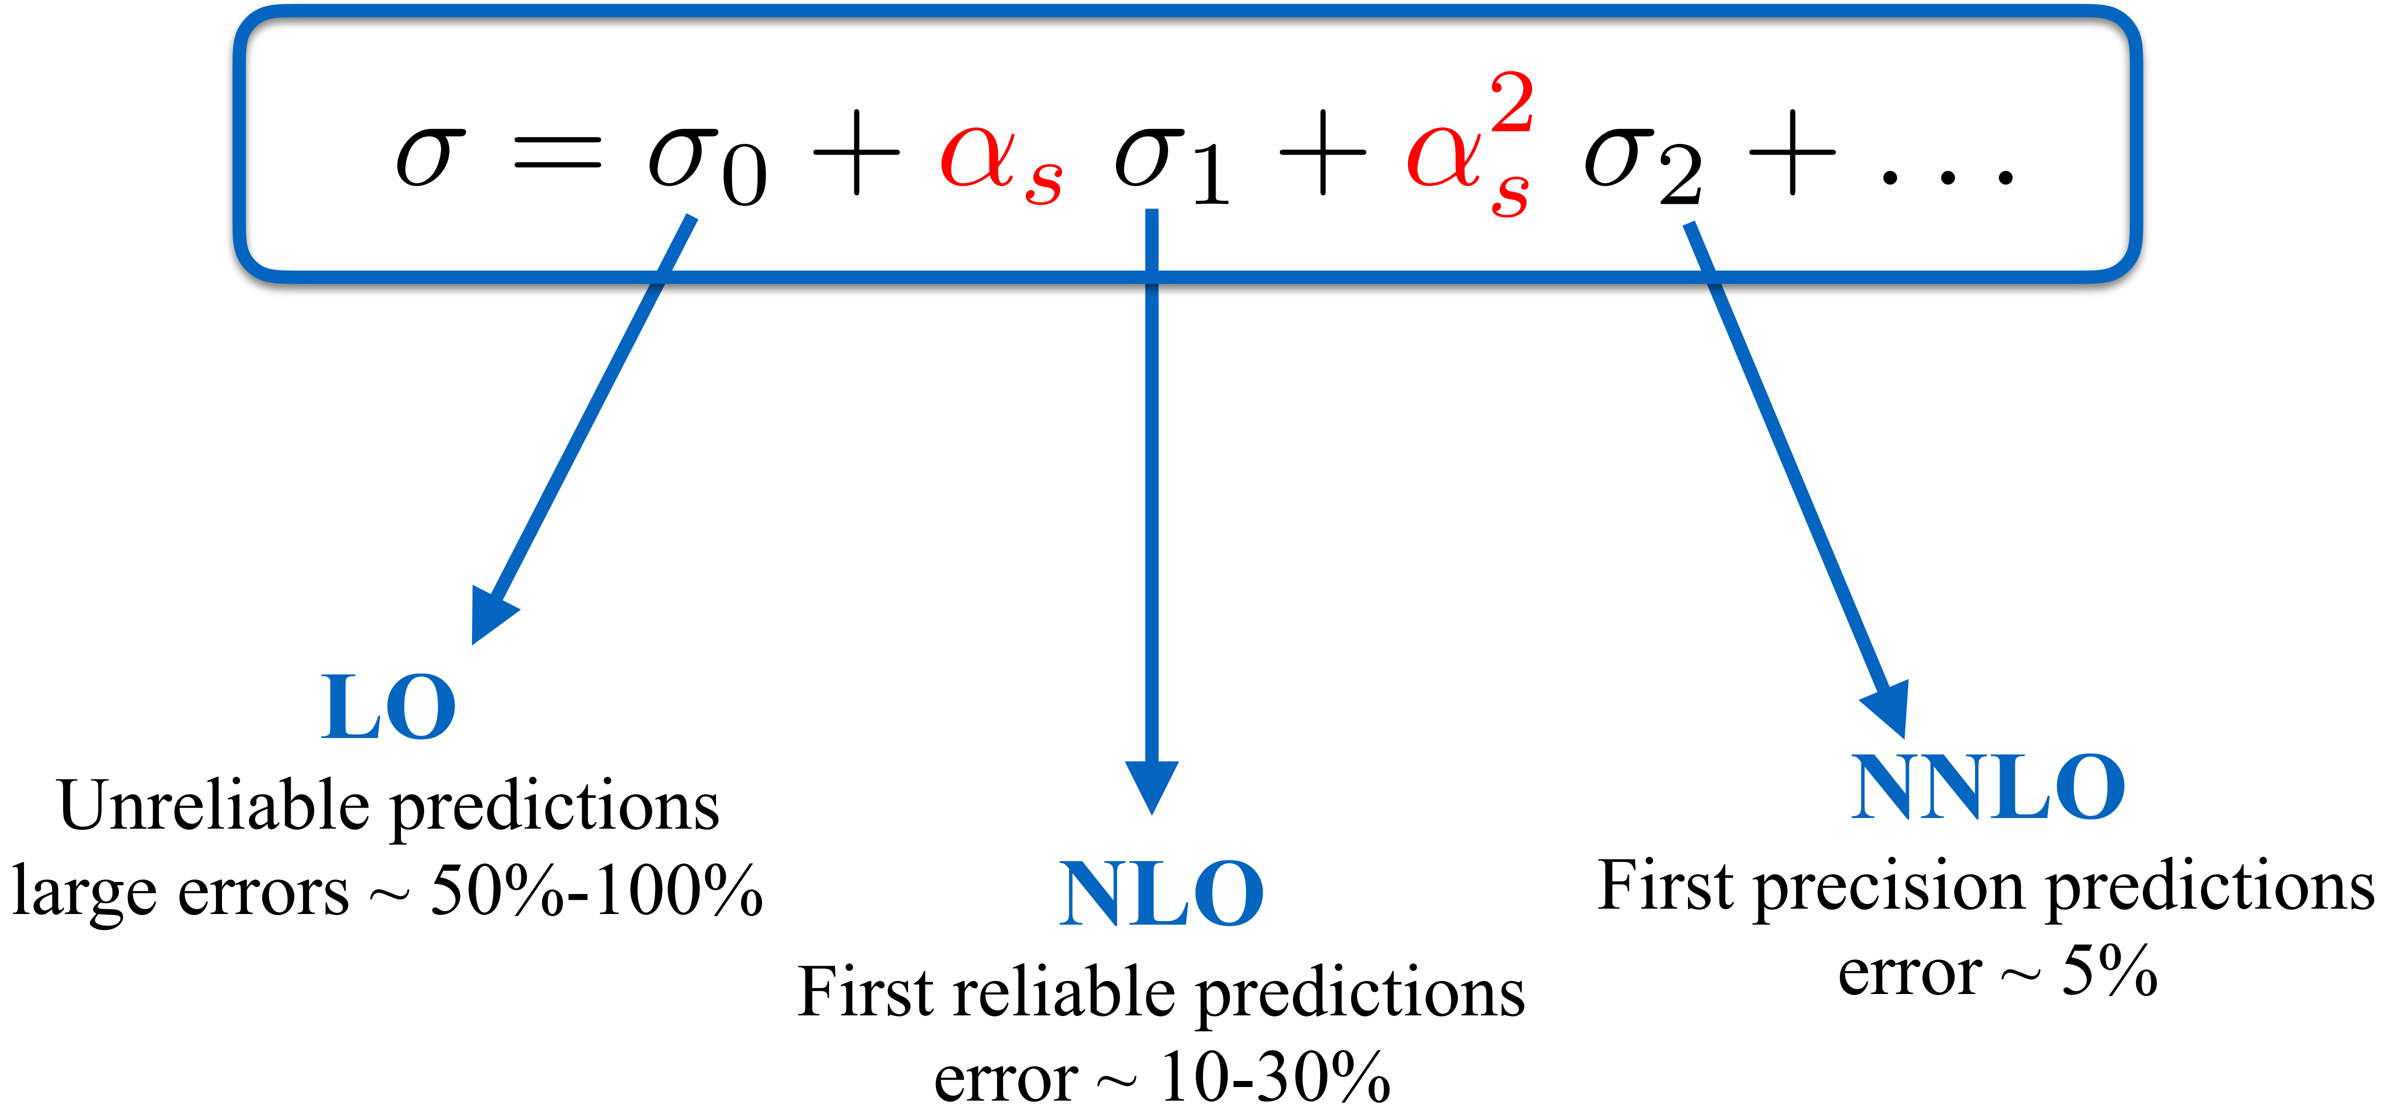

$$\boxed{\sigma = \sigma_0 + \alpha_s\,\sigma_1 + \alpha_s^2\,\sigma_2 + \ldots}$$

**LO**

Unreliable predictions
large errors ~ 50%-100%

**NLO**

First reliable predictions
error ~ 10-30%

**NNLO**

First precision predictions
error ~ 5%

# Motivations for high-precision theory

**Building for Discovery**

Strategic Plan for U.S. Particle Physics in the Global Context

From high energy physics:

*"The full discovery potential of the Higgs will be unleashed by percent-level precision studies of the Higgs properties."*

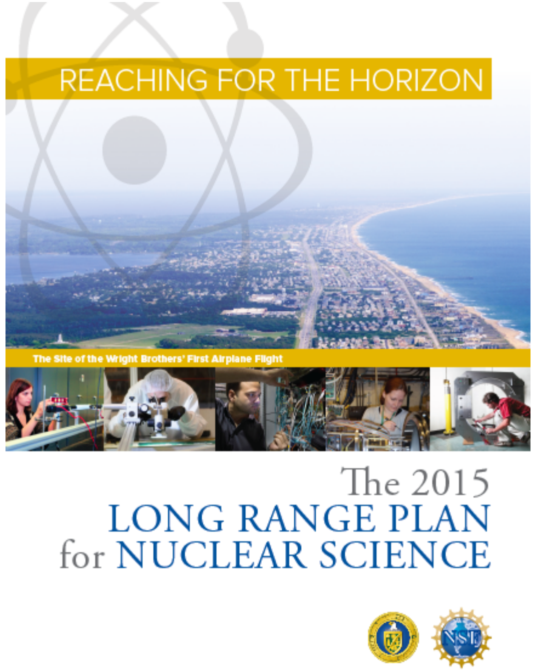

REACHING FOR THE HORIZON

The Site of the Wright Brothers' First Airplane Flight

The 2015
LONG RANGE PLAN
for NUCLEAR SCIENCE

From nuclear physics:

*"To meet the challenges and realize the full scientific potential of current and future experiments, we require new investments in theoretical and computational nuclear physics."*

# Motivations for high-precision theory

- Current data, and anticipated future data, clearly demonstrate the need for predictions beyond the current next-to-leading order (NLO) standard

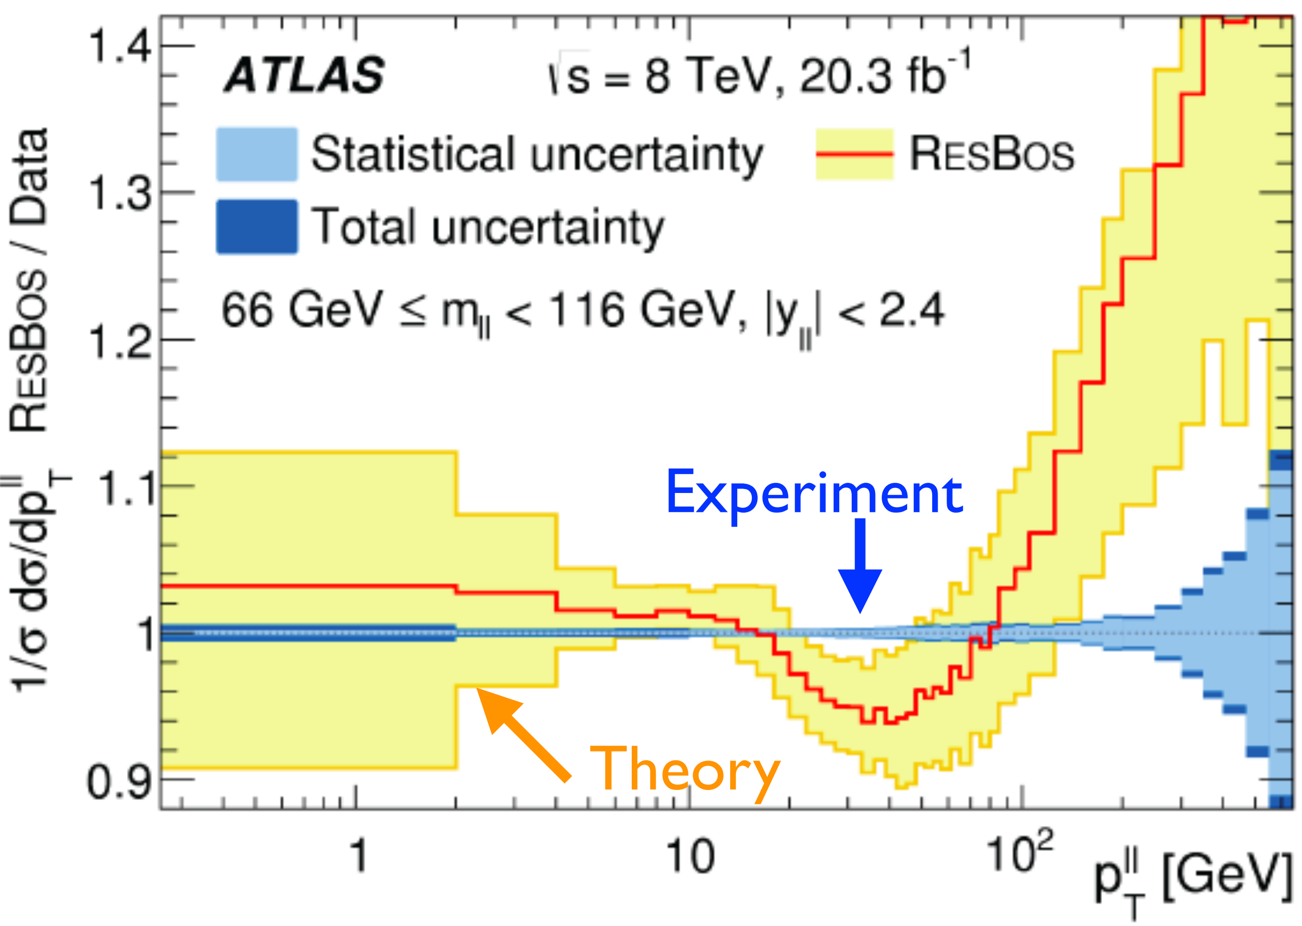

Theory errors more than an order of magnitude larger than experimental ones

Hinderer, Schlegel, Vogelsang PRD 2015

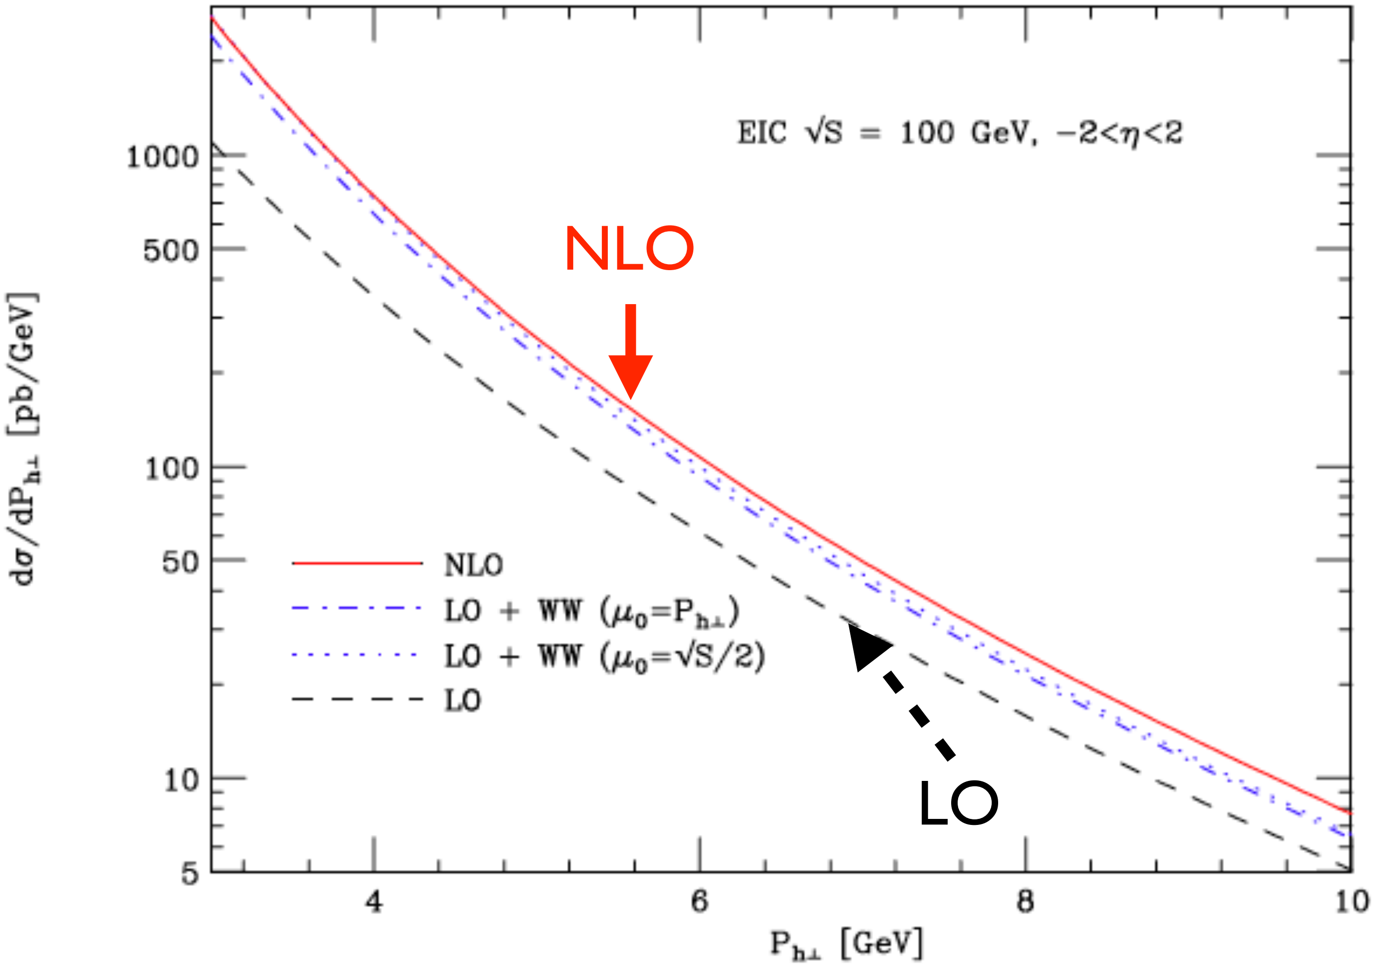

NLO corrections for EIC hadron production increase LO result by more than a factor of 2

# Anatomy of a NNLO calculation

- Next-to-next-to-leading order calculations (NNLO) in a nutshell: a collection of integrals over integrands with the following structure:

- Proliferation of possible singular configurations leads to a high-dimensional integrand with many peaks scattered throughout integration space
- Techniques require significant “by-hand” manipulation; can they be automated?
- Are the theoretical algorithms scalable to effectively use HPCs?

# An example of a traditional approach

- Sector-improved residue subtraction: an example of a more traditional approach. The integration region is divided with carefully chosen variable changes into O(200) independent sectors; each sector handles a particular configuration

Binoth, Heinrich 2000; Anastasiou, Melnikov, Petriello 2004; Czakon, 2010; RB, Melnikov, Petriello 2011

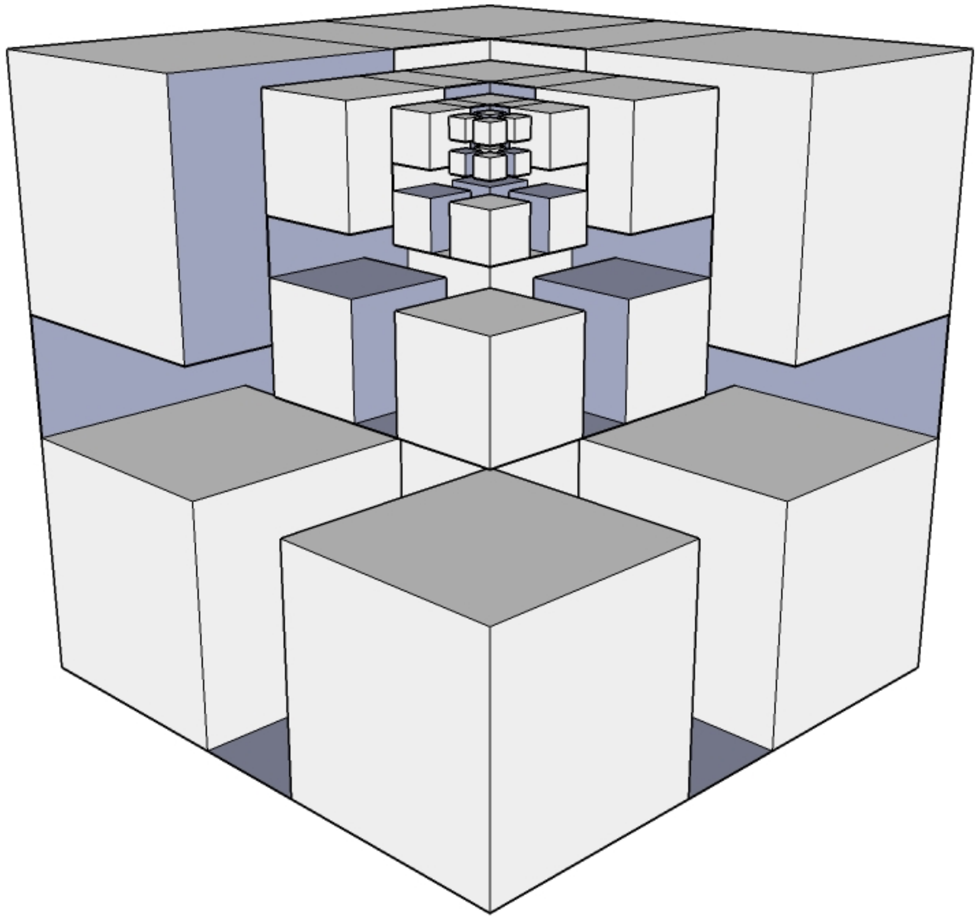

- Each sector has a different complexity and runtime; some run quickly on a single core, some require many cores to reach the requisite precision
- Significant by-hand manipulation that differs for each process

Difficult to automate and parallelize

# A new idea

## N-jettiness subtraction

RB, Focke, Liu, Petriello PRL 2015; Gaunt, Stahlhofen, Tackmann, Walsh JHEP 2015

Divide integration region into two pieces:

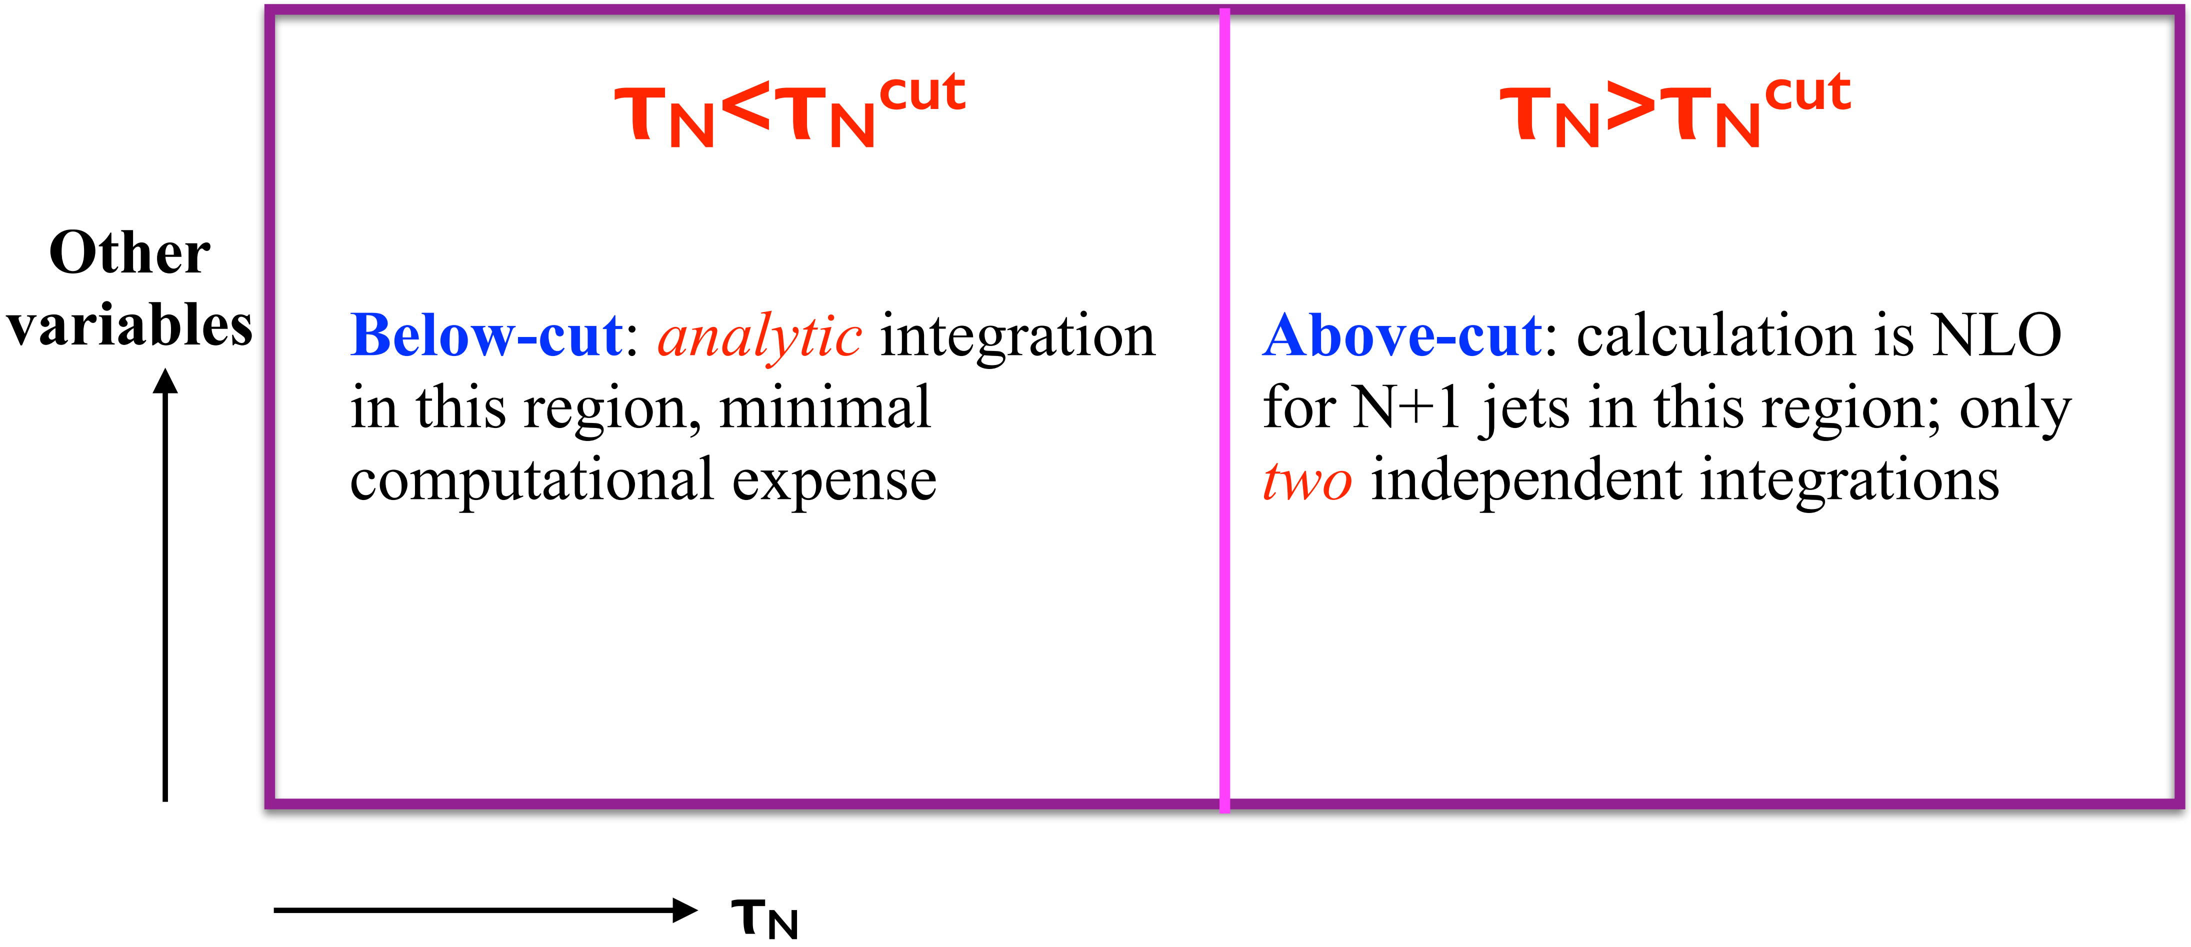

$\tau_N<\tau_N^{cut}$

$\tau_N>\tau_N^{cut}$

**Other variables**

**Below-cut**: *analytic* integration in this region, minimal computational expense

**Above-cut**: calculation is NLO for N+1 jets in this region; only *two* independent integrations

$\tau_N$

# The N-jettiness scorecard

- Advantages:

☑ Reduction of integrations: 200→3 (only 2 that have computational costs)

☑ Can reuse existing NLO frameworks for above-cut region (below-cut region is trivial to code)

☑ Much less process-dependent work, easier to release and maintain public codes incorporating many processes

- Challenges:

☑ Stringent precision requirement on the 2 non-trivial integrations (~0.01%)

(technical reason: cancellation of $\log(\tau_N^{cut})$ between above/below cut regions introduces numerical noise)

- NNLO results:

(all in past 1.5 years!)

- W+jet: **first result** (RB, Focke, Liu, Petriello PRL 2015)
- Z+jet: **first result** (RB, Campbell, Ellis, Focke, Giele, Liu, Petriello PRL 2016)
- EIC jets: **first result** (Abelof, RB, Liu, Petriello, 2016 )
- Higgs+jet (RB, Focke, Giele, Liu, Petriello PLB 2015)
- Higgs,W,Z,ZH,WH (RB, Campbell, Ellis, Focke, Giele, Liu, Petriello, Williams 2016, MCFM 8.0)
- $\gamma\gamma$: corrected previous result in literature (Campbell, Ellis, Li, Williams, 2016)

Use of HPC instrumental in addressing this challenge!

# Overview of HPC implementation

- Use adaptive Monte Carlo integration (VEGAS) for integrations
- Fundamental assumption: evaluation of integrand dominates runtime

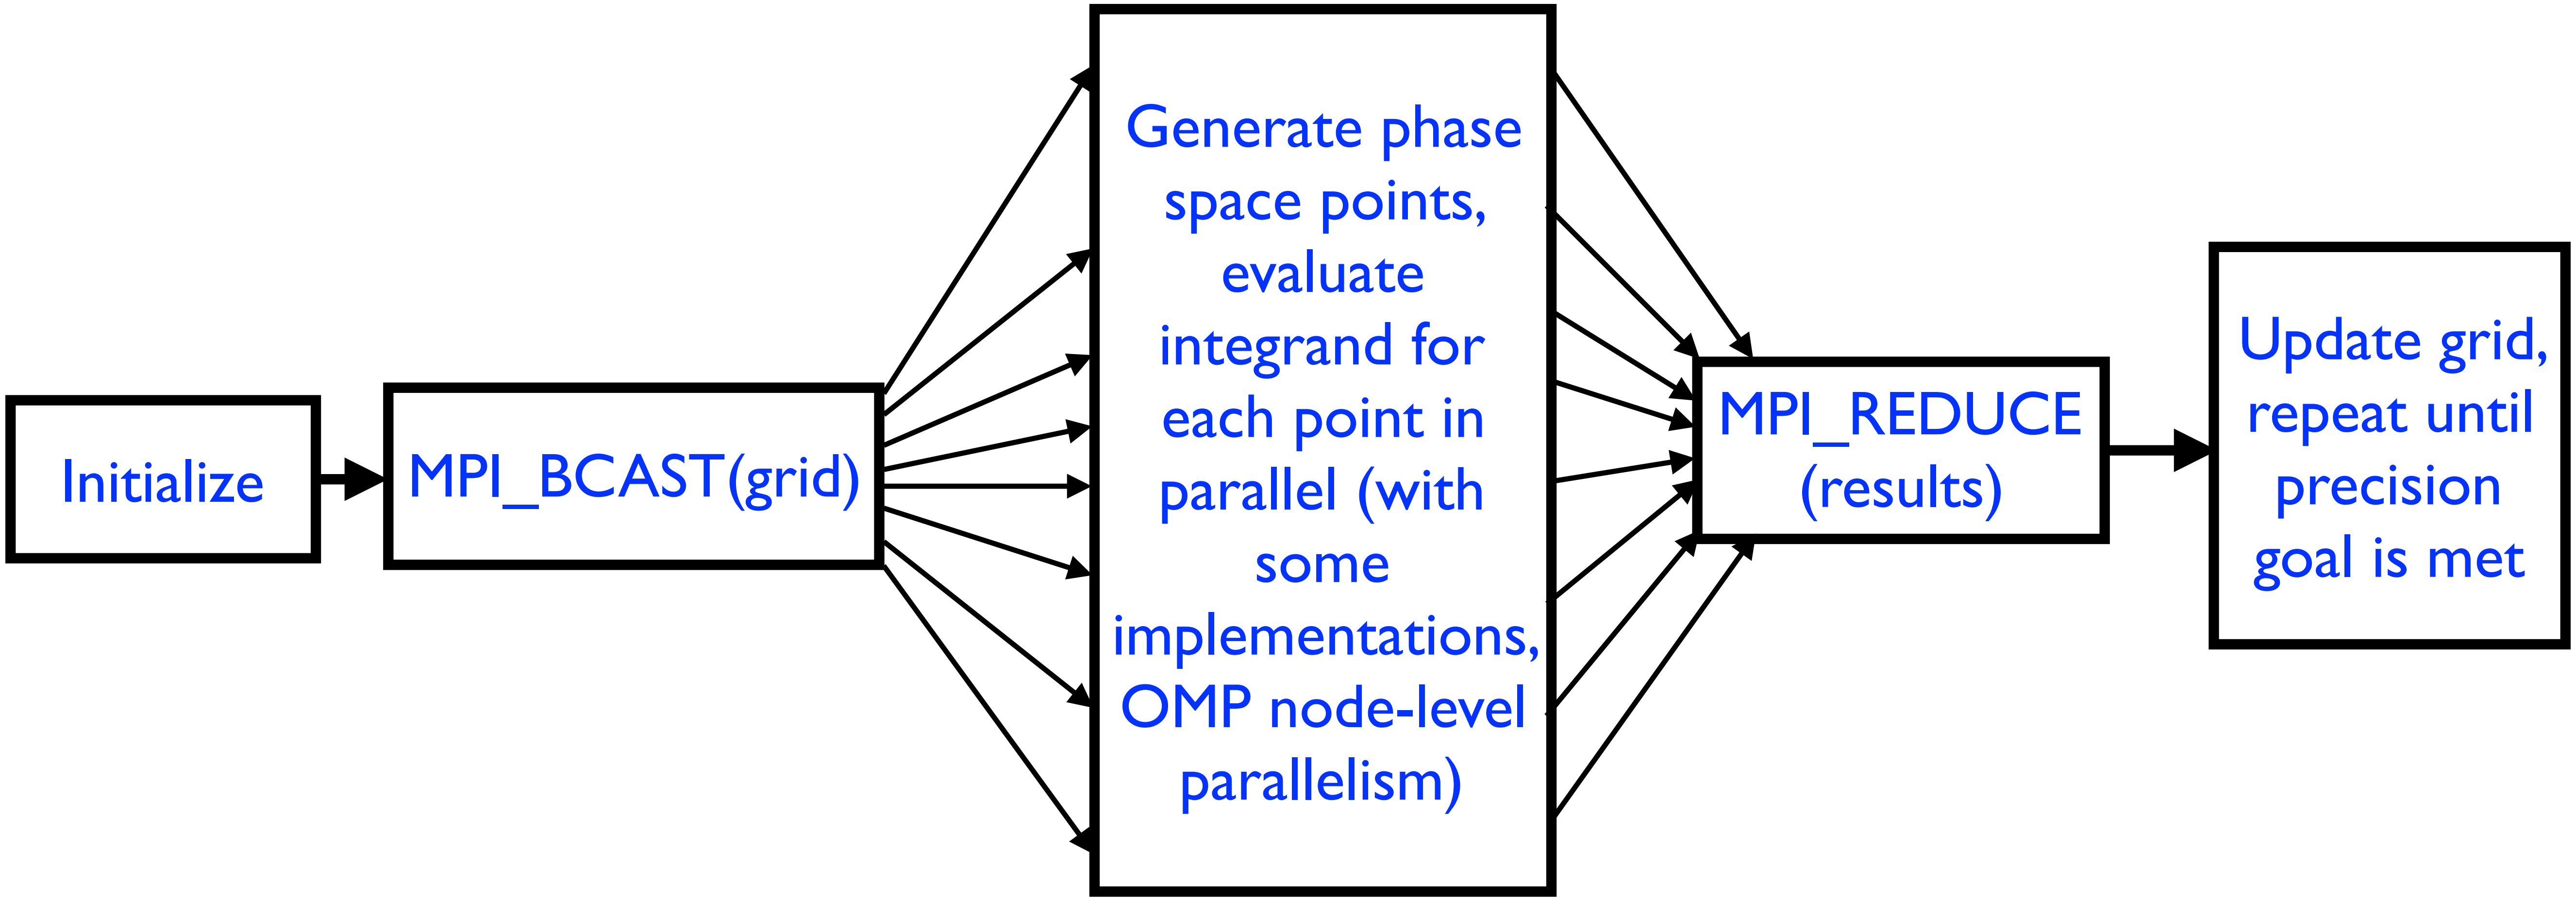

Initialize

MPI_BCAST(grid)

Generate phase space points, evaluate integrand for each point in parallel (with some implementations, OMP node-level parallelism)

MPI_REDUCE (results)

Update grid, repeat until precision goal is met

- Performance measure: strong scaling of the implementation

# Mira scaling study

- Here is a representative example: scaling study of the W+jet@NNLO process on Mira (similar results for other processes, and on Edison@NERSC)

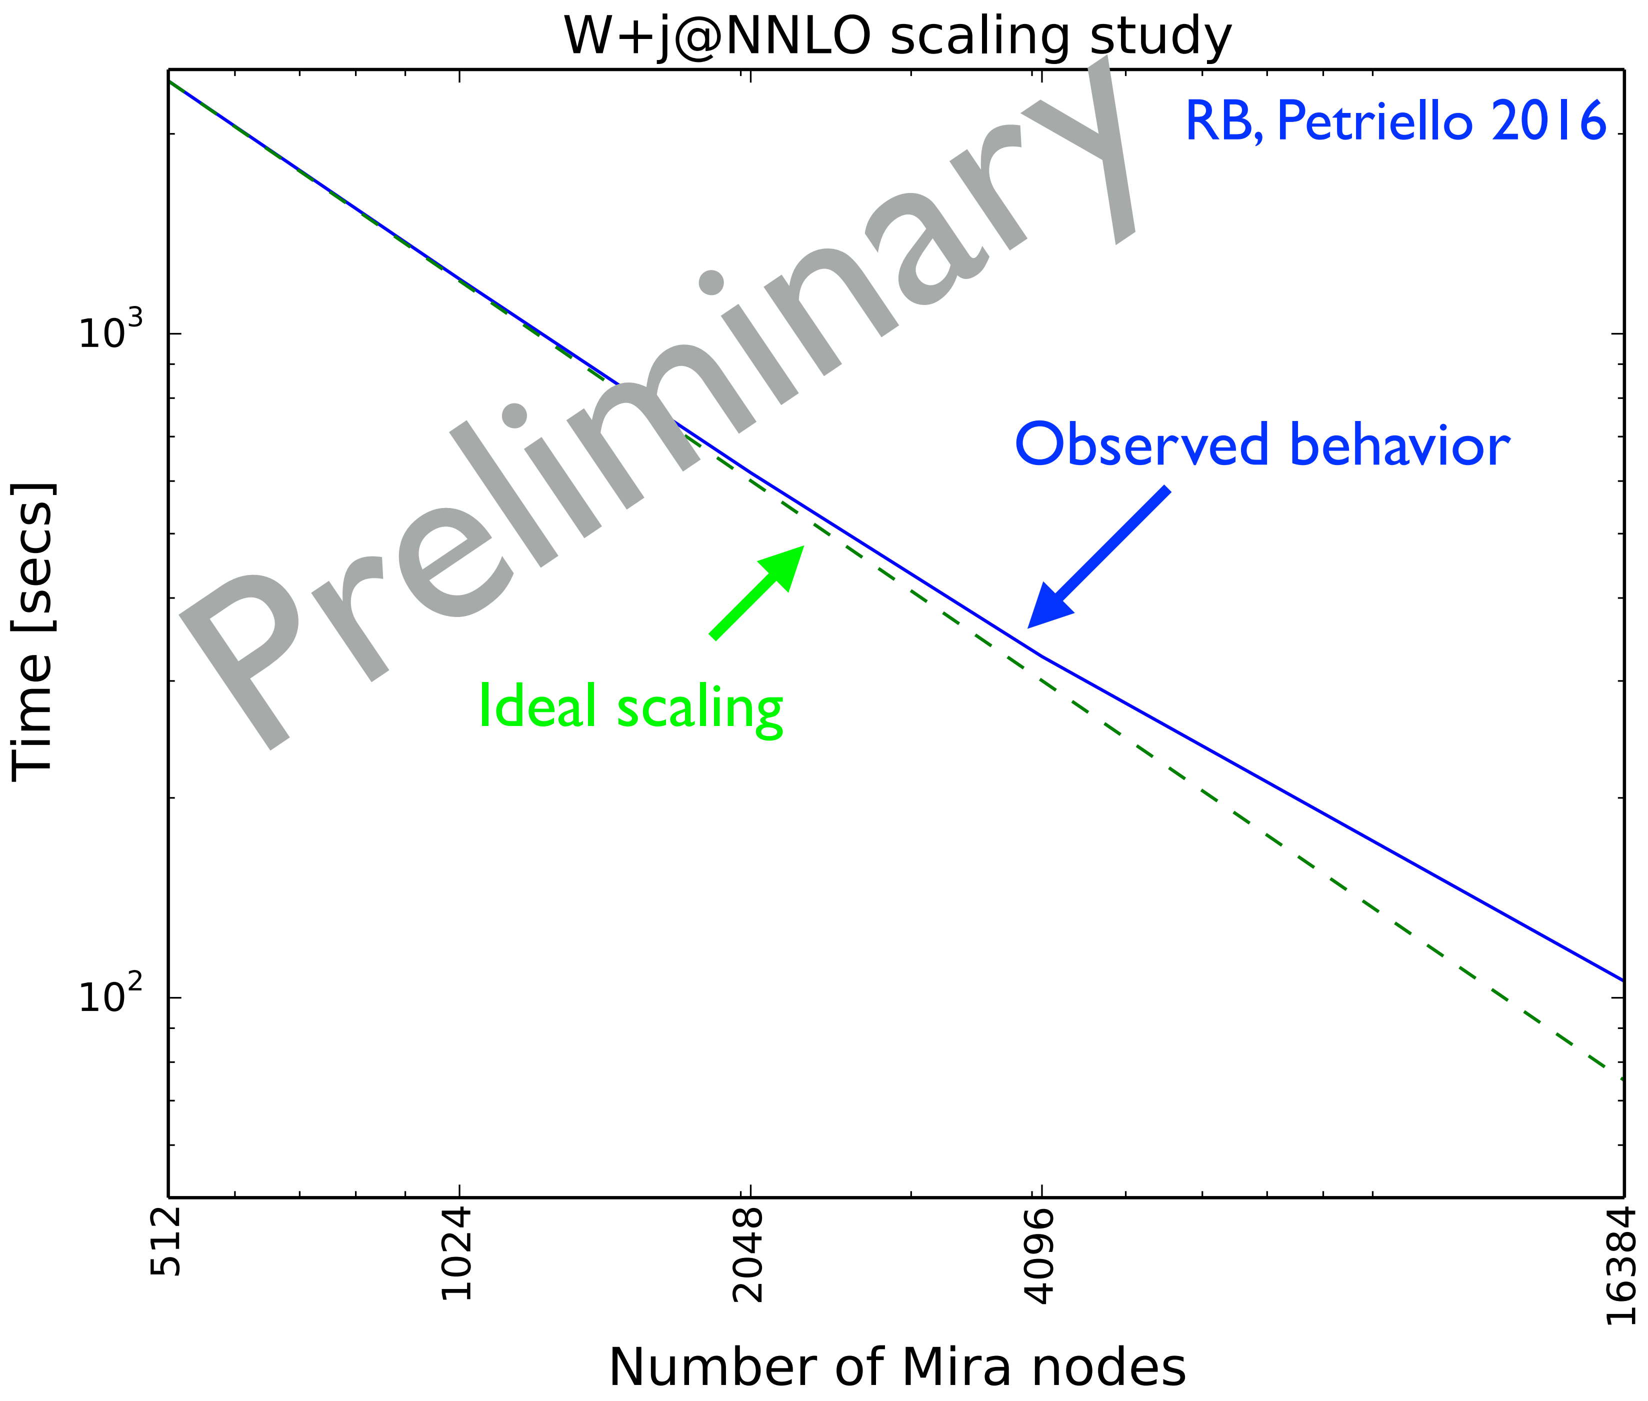

- 20% deviation from ideal scaling for a 16K Mira partition
- Known cause: VEGAS grid adaptation on a single core
- Can be improved, but we usually run single integrations on 4K or less partitions
- Future: scaling improves as we tackle more complicated processes!

# N-jettiness implementations

- **MCFM 8.0** (RB, Campbell, Ellis, Focke, Giele, Liu, Petriello, Williams 2016)
  - Implementation of color-singlet processes: H, W, Z, WH, ZH, γγ
  - Fortran77+Fortran90
  - MPI+OpenMP VEGAS algorithm
  - Ongoing development for color singlet+jet processes
- **DISTRESS** (Abelof, RB, Liu, Petriello 2016)
  - Implementation of jet production in deeply-inelastic scattering (DIS)
  - Fortran 90
  - MPI VEGAS algorithm (OMP under development)
  - Ongoing development for other DIS processes

# HPC resources

- Director’s Discretionary QCDHPC (PI: RB)
- ALCC EnergyFEC (PIs: Childers, LeCompte)

- m2405: QCDNNLO (PI: RB)
- m1758: THEP-HPC (PI: Hoeche)

# HPC resources

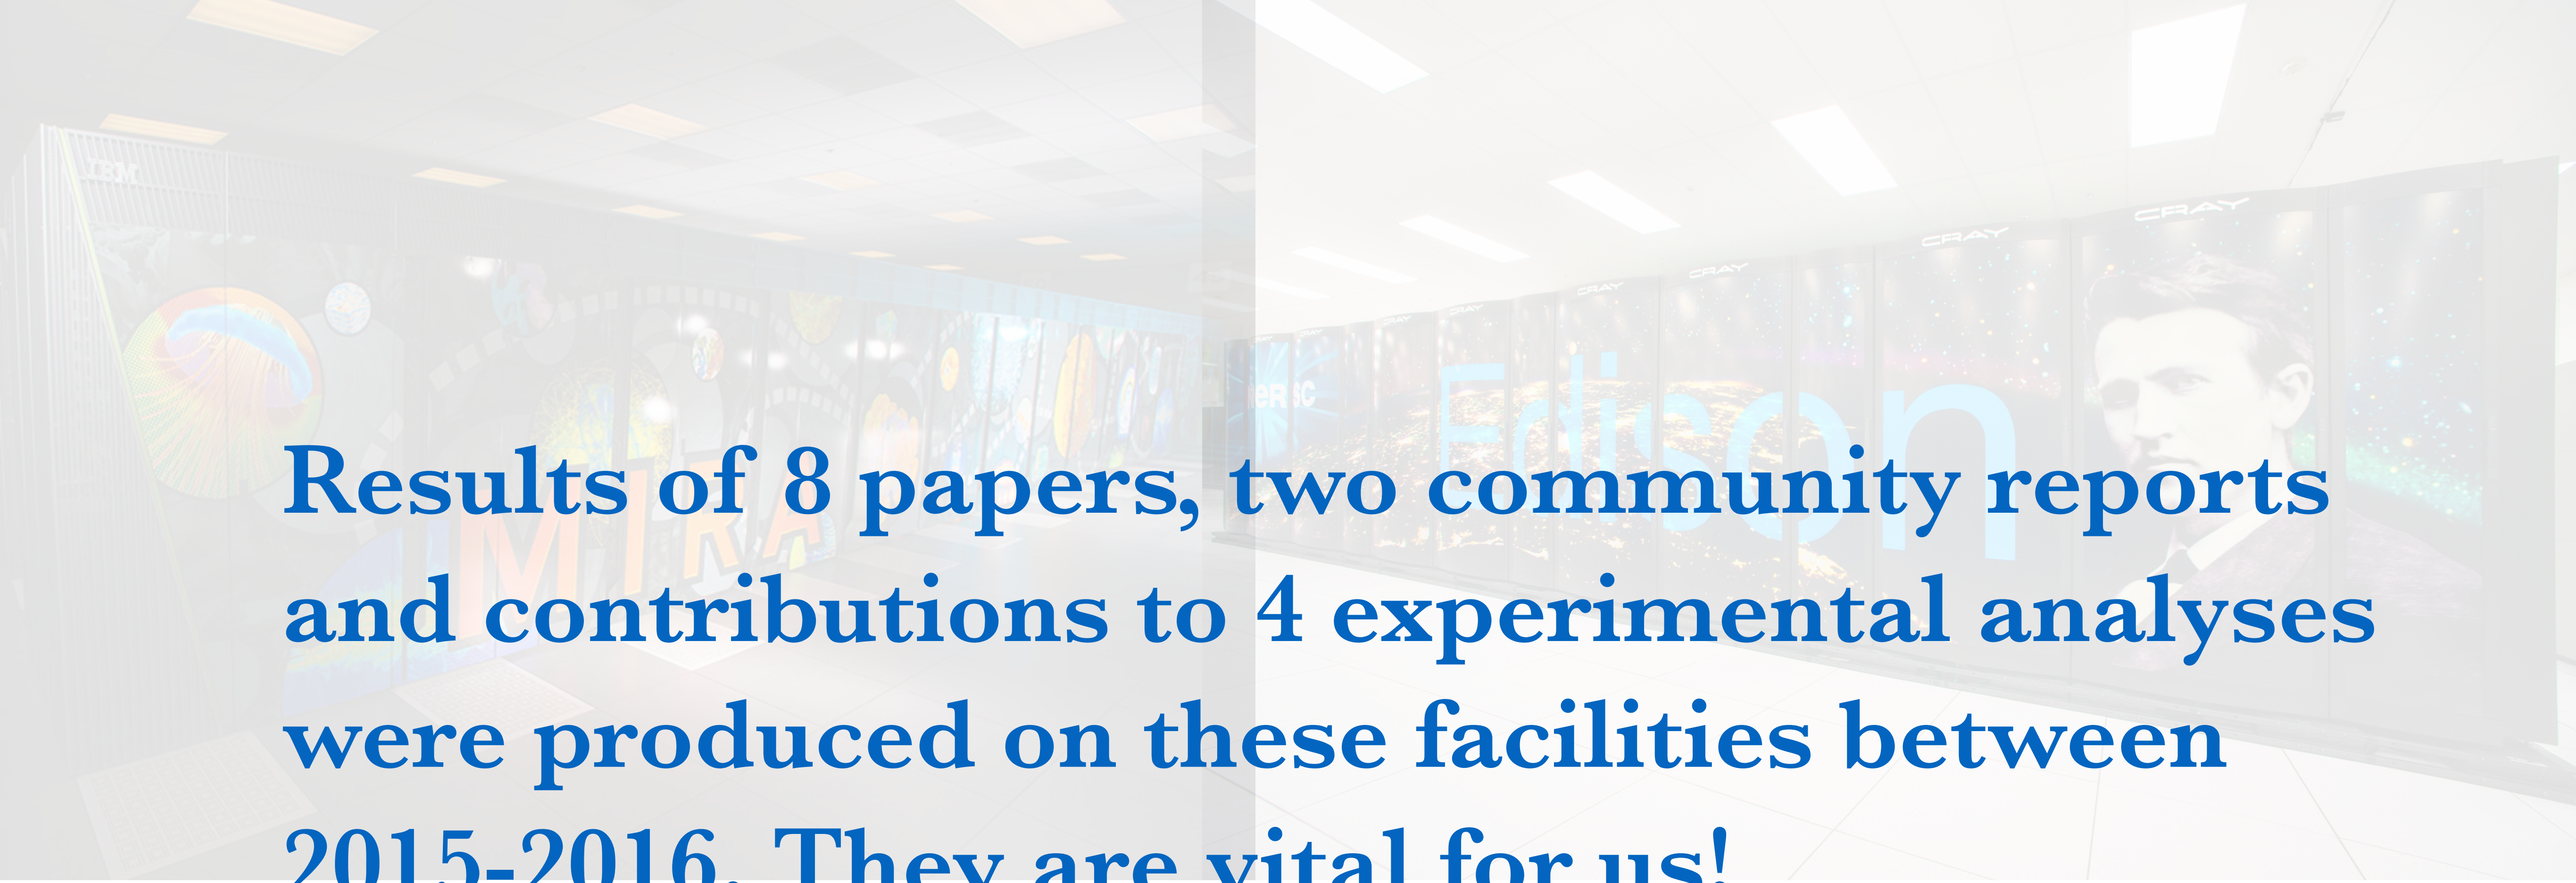

**Results of 8 papers, two community reports and contributions to 4 experimental analyses were produced on these facilities between 2015-2016. They are vital for us!**

- Director’s Discretionary QCDHPC (PI: RB)
- ALCC EnergyFEC (PIs: Childers, LeCompte)
- m2405: Ultra-precise Predictions for the Terascale (PI: RB)
- m1758: HPC for QCD Theory (PI: Hoeche)

# Details of running on Mira

- Typically run for several $\tau_N^{cut}$ values to monitor systematic effects on the integration; also for several different seeds for the random number generator
- All jobs should finish at the same time to allow for systematic error study, statistical averaging, and analysis (all done off-site with in-house python scripts)
- Sweet spot: single-grid integrations on 1024/2048-node partitions, bundled into 16K ensemble jobs (may change for future applications!)

Example single job:

```
qsub -A EnergyFEC -t 45 -n 1024 --env I_MPI_PLATFORM=auto
--mode c64 ./distress y0=1d-14 tau1cut=0.00005 seed=800
outfile=tau00005_s800.dat
```

Example ensemble job:

```
qsub -A EnergyFEC -t 60 -n 16384 --disable_preboot
--mode script enscript_etajet.sh
```

# Details of running on Mira

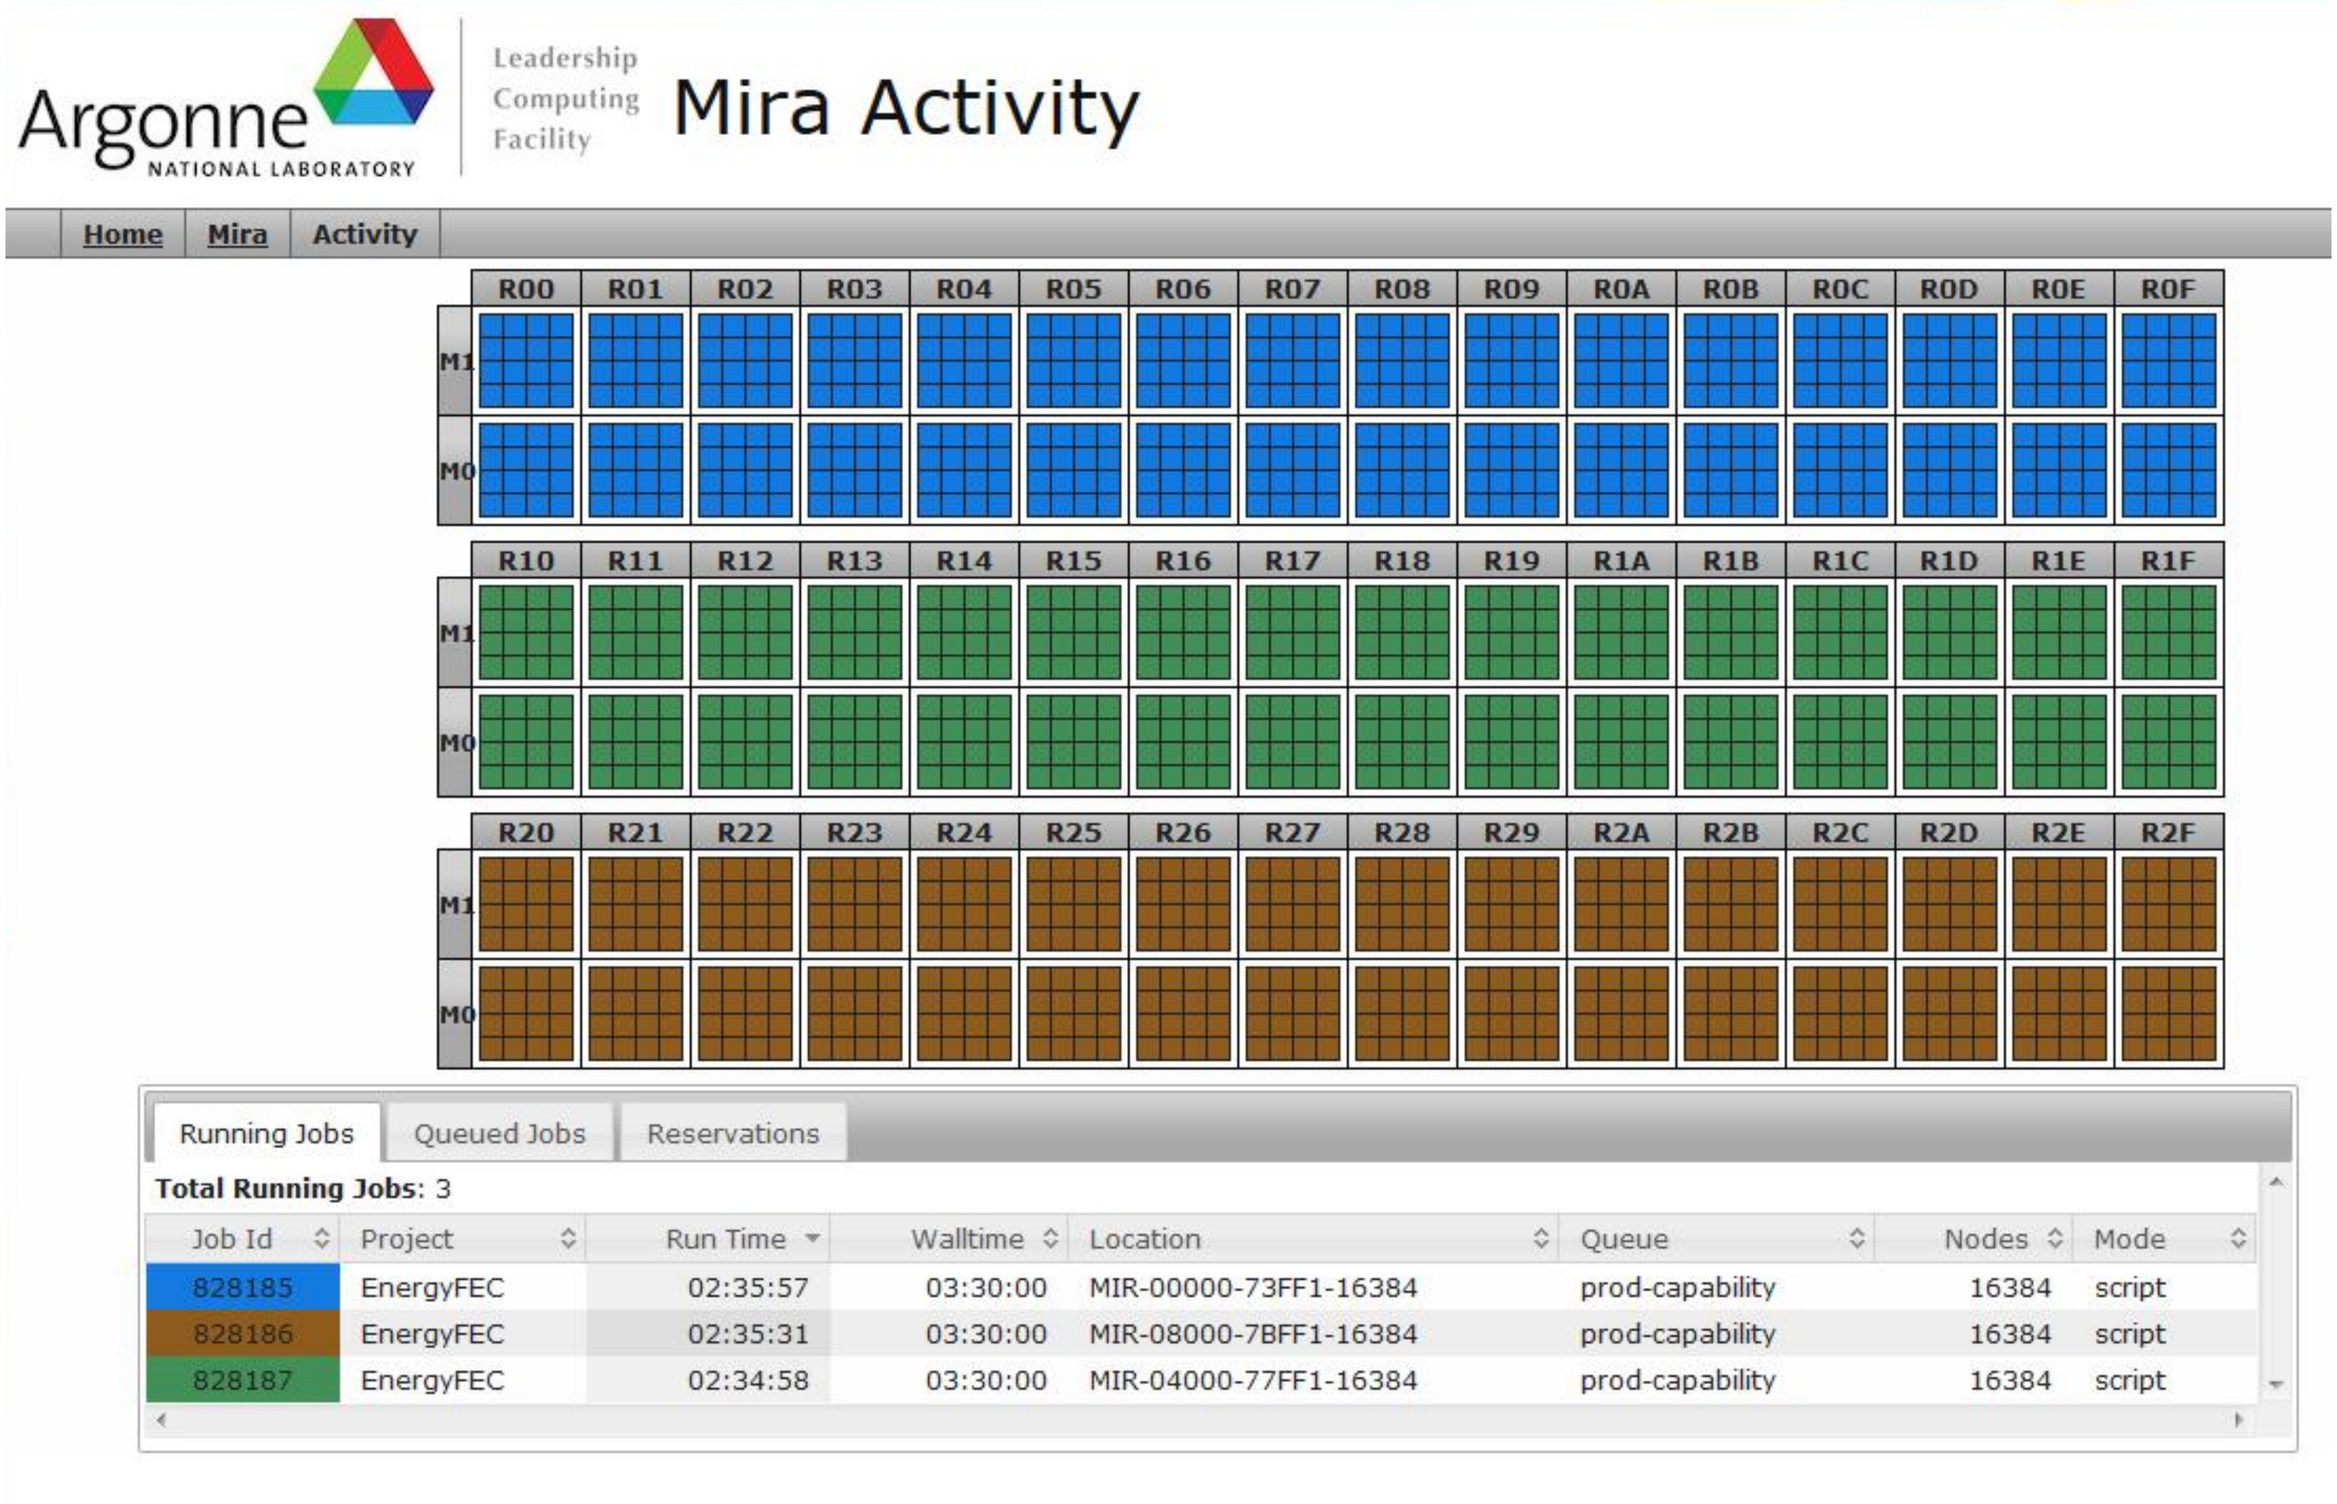

| | R00 | R01 | R02 | R03 | R04 | R05 | R06 | R07 | R08 | R09 | R0A | R0B | R0C | R0D | R0E | R0F |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| M1 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 |
| M0 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 | 🟦 |

| | R10 | R11 | R12 | R13 | R14 | R15 | R16 | R17 | R18 | R19 | R1A | R1B | R1C | R1D | R1E | R1F |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| M1 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 |
| M0 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 | 🟩 |

| | R20 | R21 | R22 | R23 | R24 | R25 | R26 | R27 | R28 | R29 | R2A | R2B | R2C | R2D | R2E | R2F |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| M1 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 |
| M0 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 | 🟫 |

**Total Running Jobs**: 3

| Job Id | Project | Run Time | Walltime | Location | Queue | Nodes | Mode |
|---|---|---|---|---|---|---|---|
| 🟦 828185 | EnergyFEC | 02:35:57 | 03:30:00 | MIR-00000-73FF1-16384 | prod-capability | 16384 | script |
| 🟫 828186 | EnergyFEC | 02:35:31 | 03:30:00 | MIR-08000-7BFF1-16384 | prod-capability | 16384 | script |
| 🟩 828187 | EnergyFEC | 02:34:58 | 03:30:00 | MIR-04000-77FF1-16384 | prod-capability | 16384 | script |

We will happily use the entire machine!

# Details of running on Mira

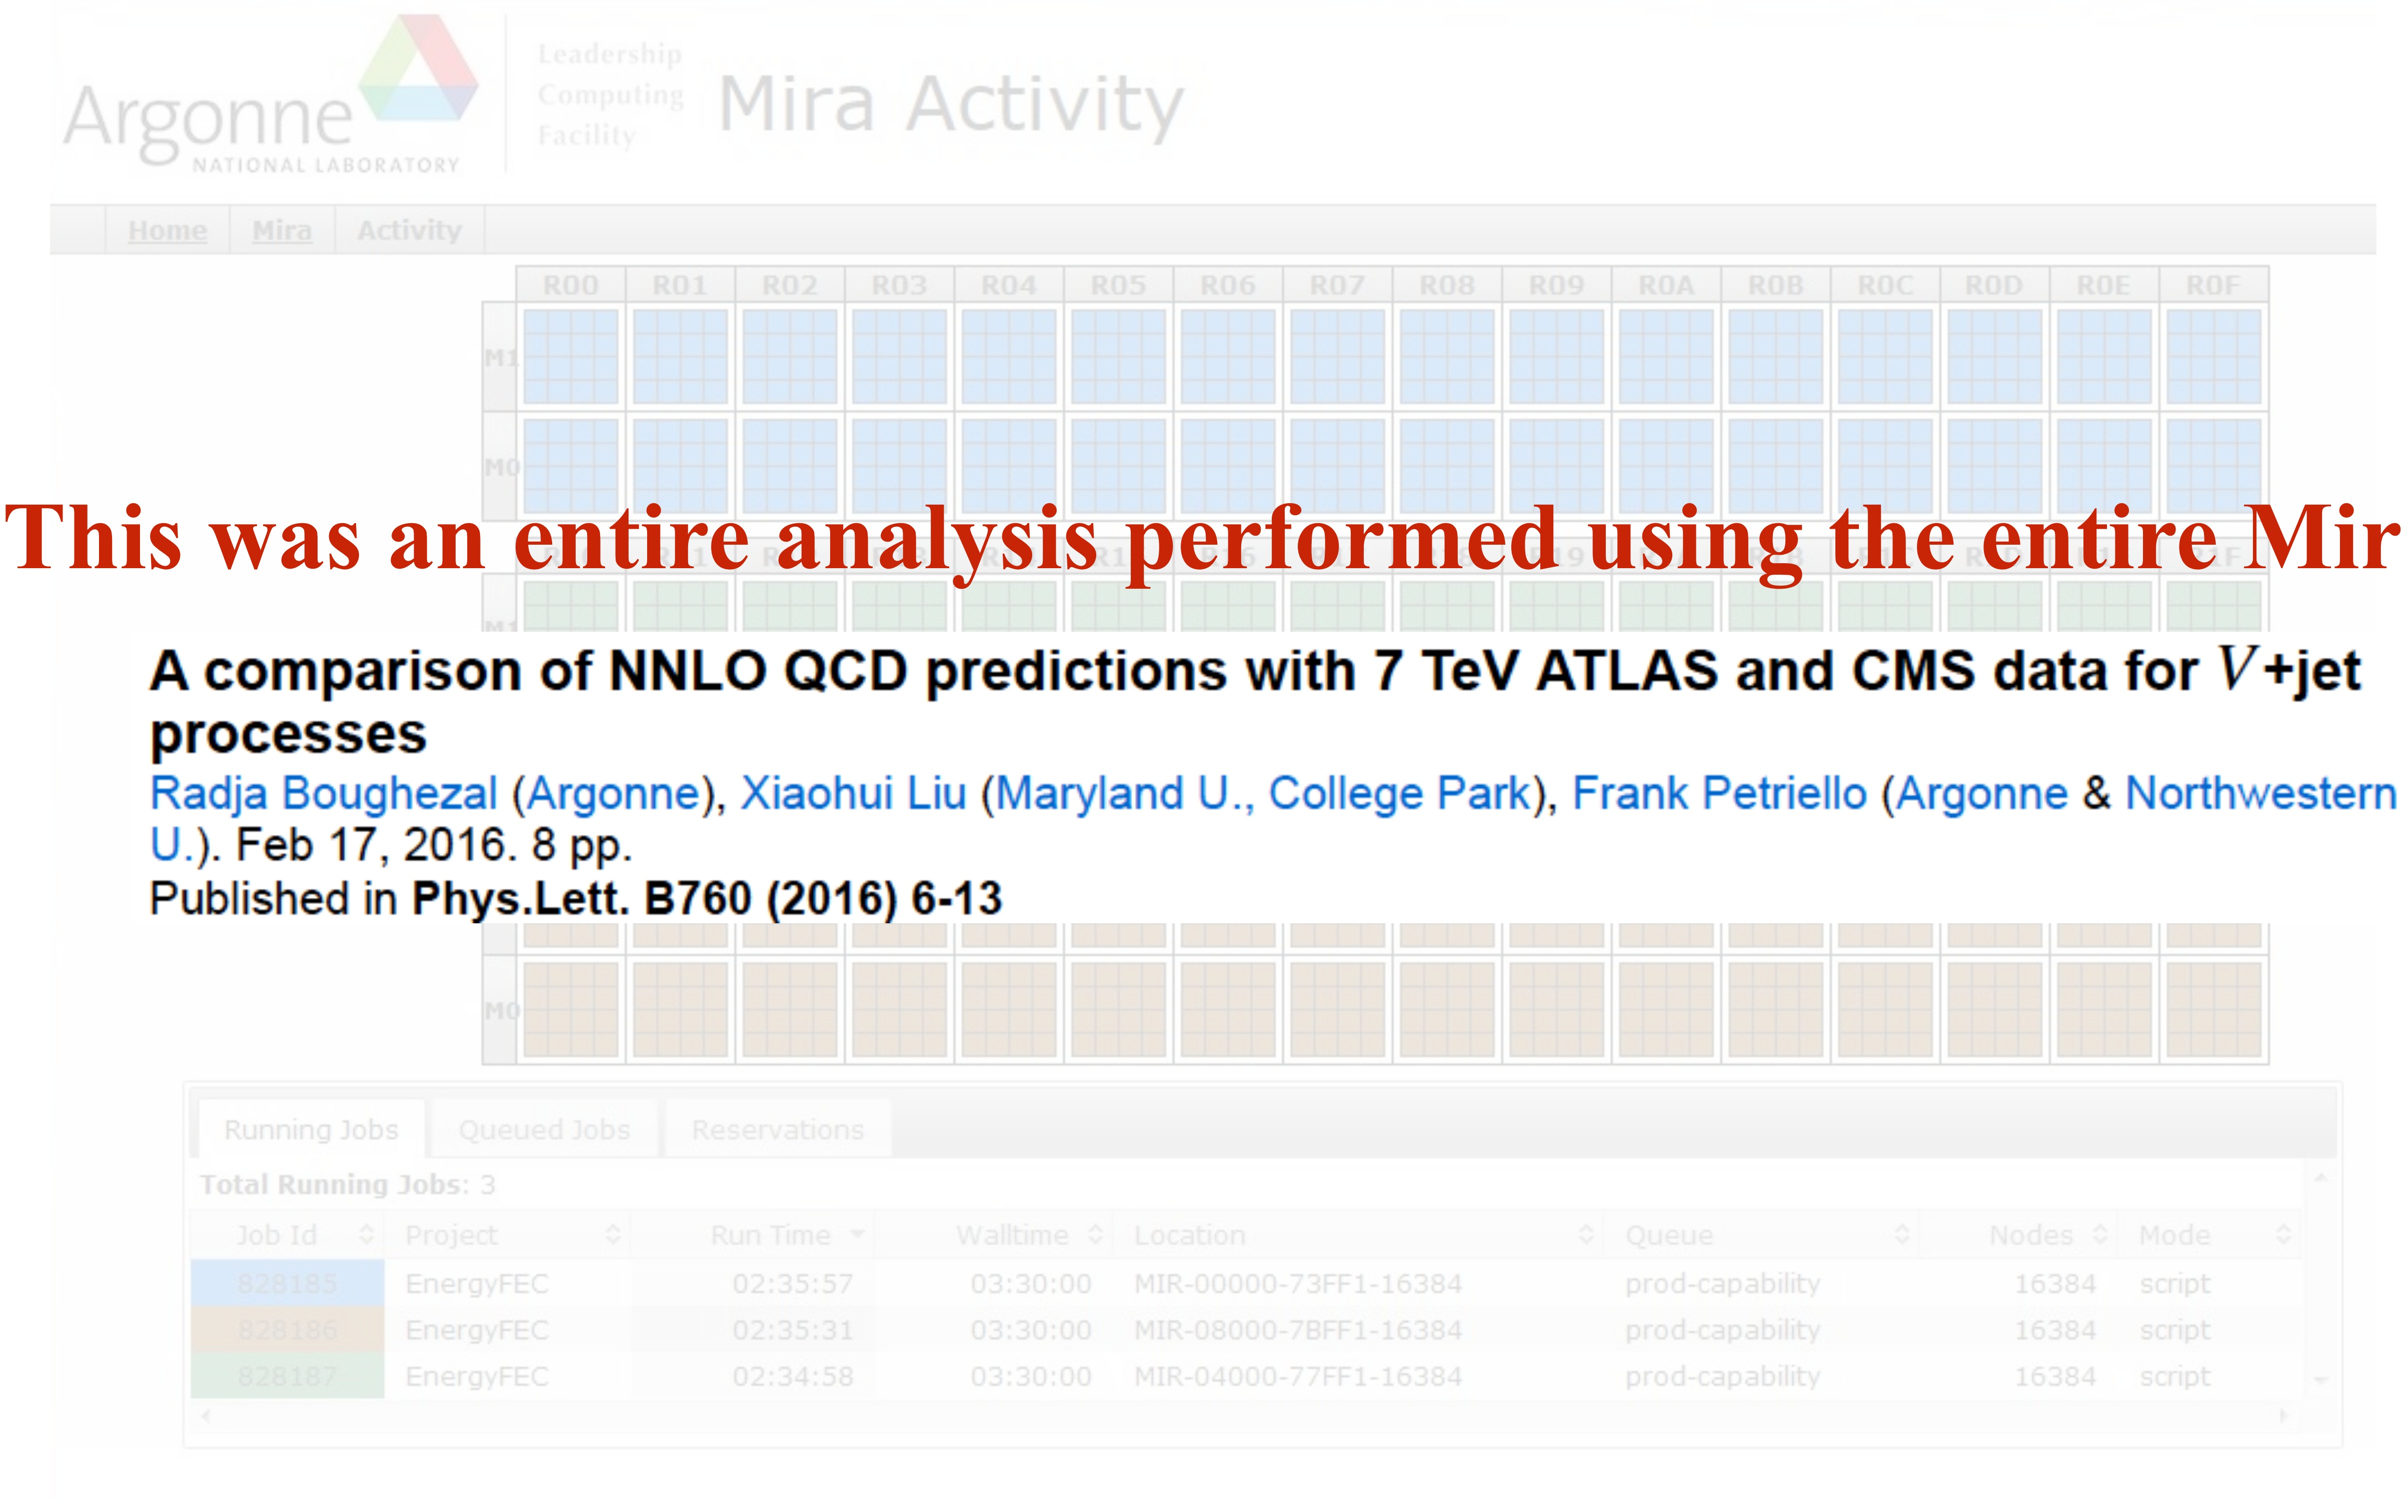

Total Running Jobs: 3

| Job Id | Project | Run Time | Walltime | Location | Queue | Nodes | Mode |
|---|---|---|---|---|---|---|---|
| 828185 | EnergyFEC | 02:35:57 | 03:30:00 | MIR-00000-73FF1-16384 | prod-capability | 16384 | script |
| 828186 | EnergyFEC | 02:35:31 | 03:30:00 | MIR-08000-7BFF1-16384 | prod-capability | 16384 | script |
| 828187 | EnergyFEC | 02:34:58 | 03:30:00 | MIR-04000-77FF1-16384 | prod-capability | 16384 | script |

**This was an entire analysis performed using the entire Mira!**

**A comparison of NNLO QCD predictions with 7 TeV ATLAS and CMS data for $V$+jet processes**

Radja Boughezal (Argonne), Xiaohui Liu (Maryland U., College Park), Frank Petriello (Argonne & Northwestern U.). Feb 17, 2016. 8 pp.

Published in **Phys.Lett. B760 (2016) 6-13**

We will happily use the entire machine!

# Details of running on Edison

- Same run parameter considerations as on Mira (multiple $\tau_N^{cut}$, VEGAS seeds)
- Observe good scaling with DISTRESS: run 684-node jobs to benefit from reduced charge factor (1.2 versus 2.0)

Example single job submission script:

```
#!/bin/bash -l
#SBATCH -p regular
#SBATCH -N 684
#SBATCH -A m2405
#SBATCH -t 01:45:00
#SBATCH -J tau00001_s7800_R

cd $SLURM_SUBMIT_DIR

srun -n 16416 ./distress y0=1d-15 tau1cut=0.00001 seed=7800
outfile='tau00001_s7800_R.dat'
```

# Validation of results

- A powerful check of the formalism is the independence of the final result from $\tau_N^{cut}$, which checks the implementation of almost all parts of the calculation

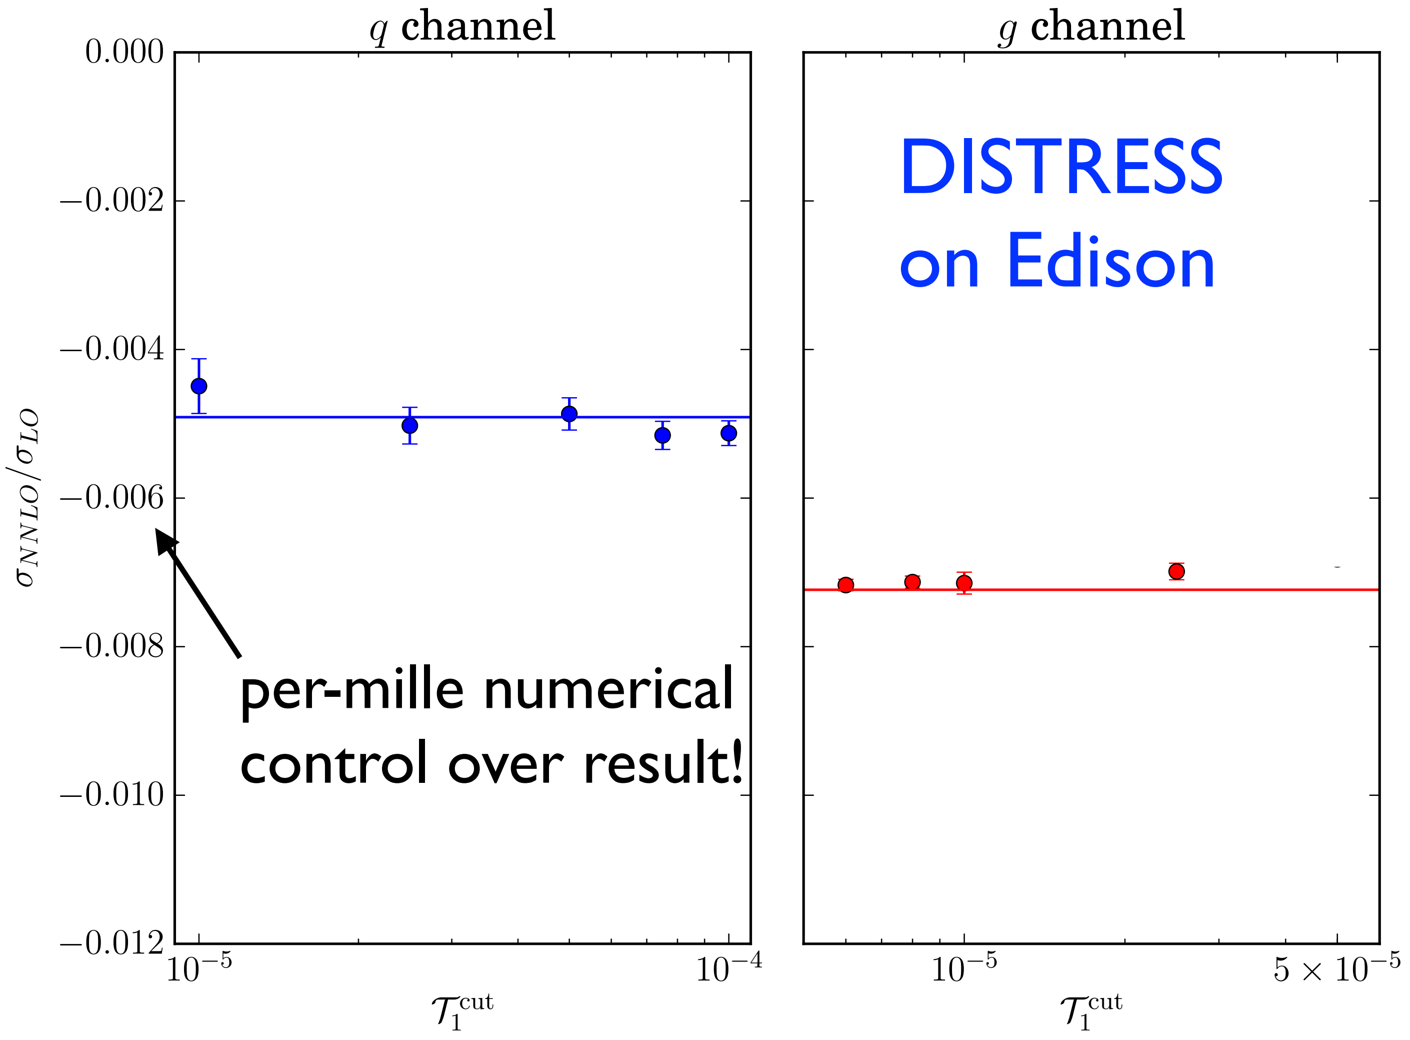

DISTRESS on Edison

per-mille numerical control over result!

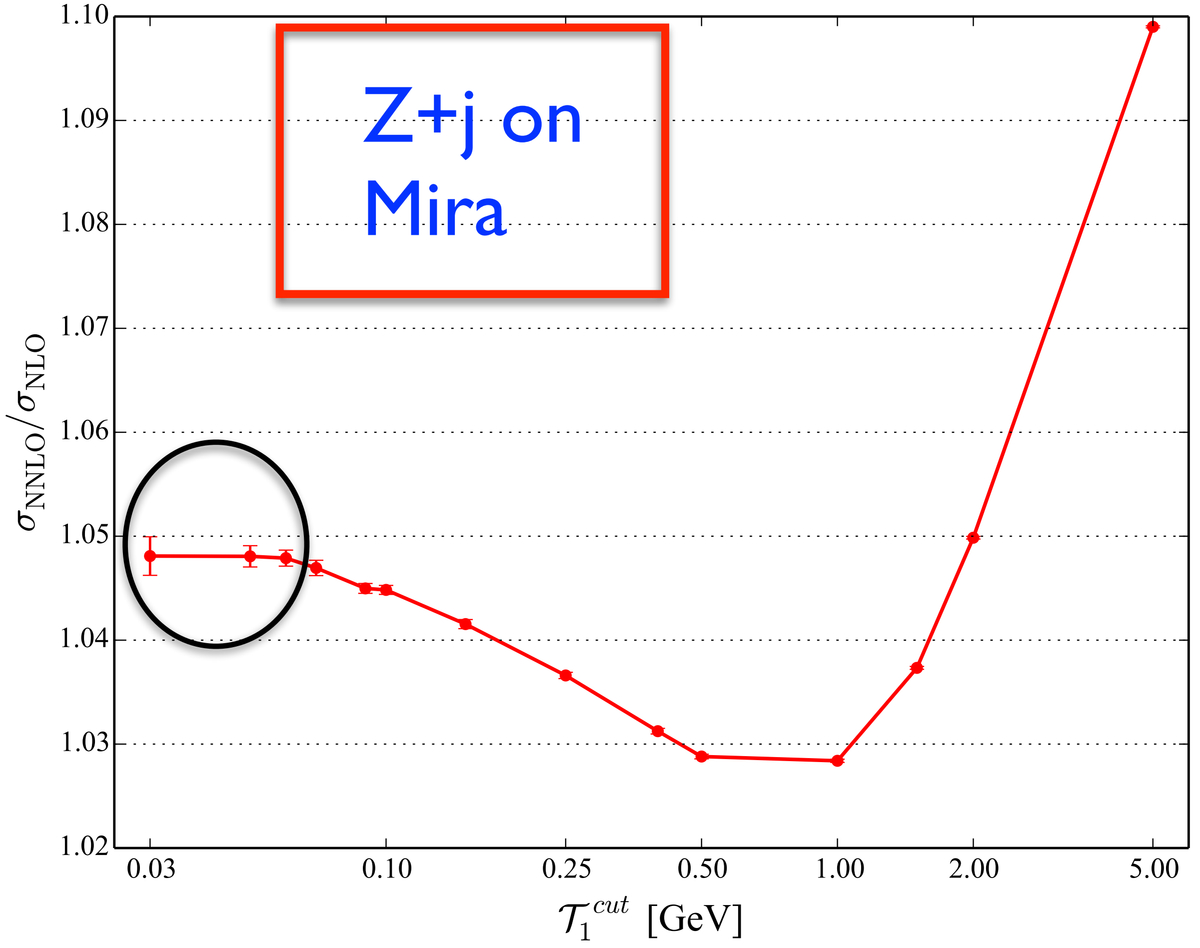

Z+j on Mira

- solid lines: inclusive structure functions
- Obtained from DISTRESS after integration over phase space, also known analytically

- Result taken from circled region, where result becomes independent of the cut

# Scientific highlights

- Runs on Mira for the first time satisfactorily describe data for Z+jet coming from ATLAS and CMS

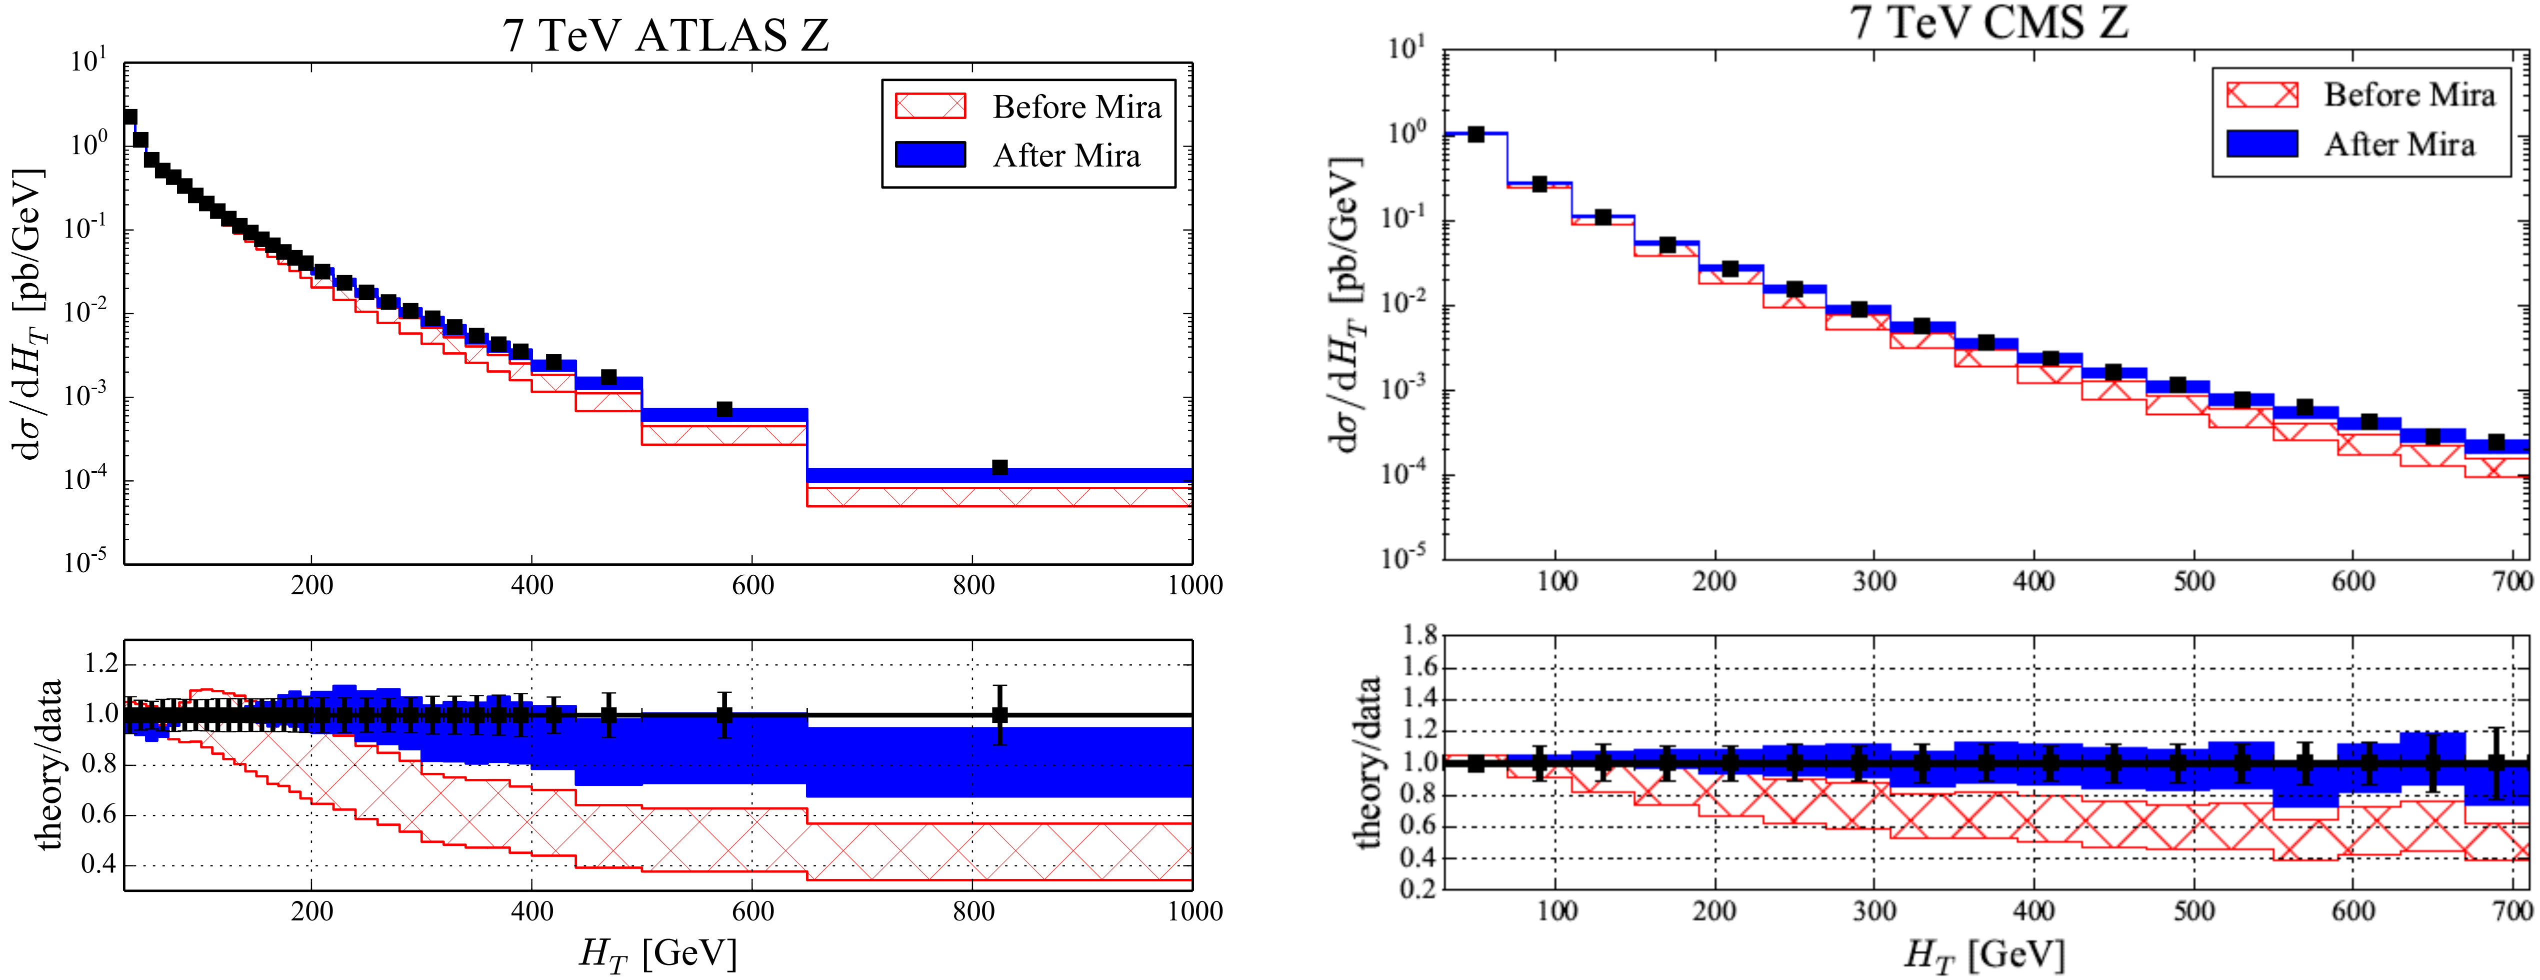

RB, Liu, Petriello PLB 2016

Argonne NATIONAL LABORATORY | Leadership Computing Facility

# Scientific highlights

- Large corrections observed for jet production at an Electron Ion collider tamed at NNLO
- DISTRESS allows for a detailed investigation of how the different parts of the corrections combine to give the full result

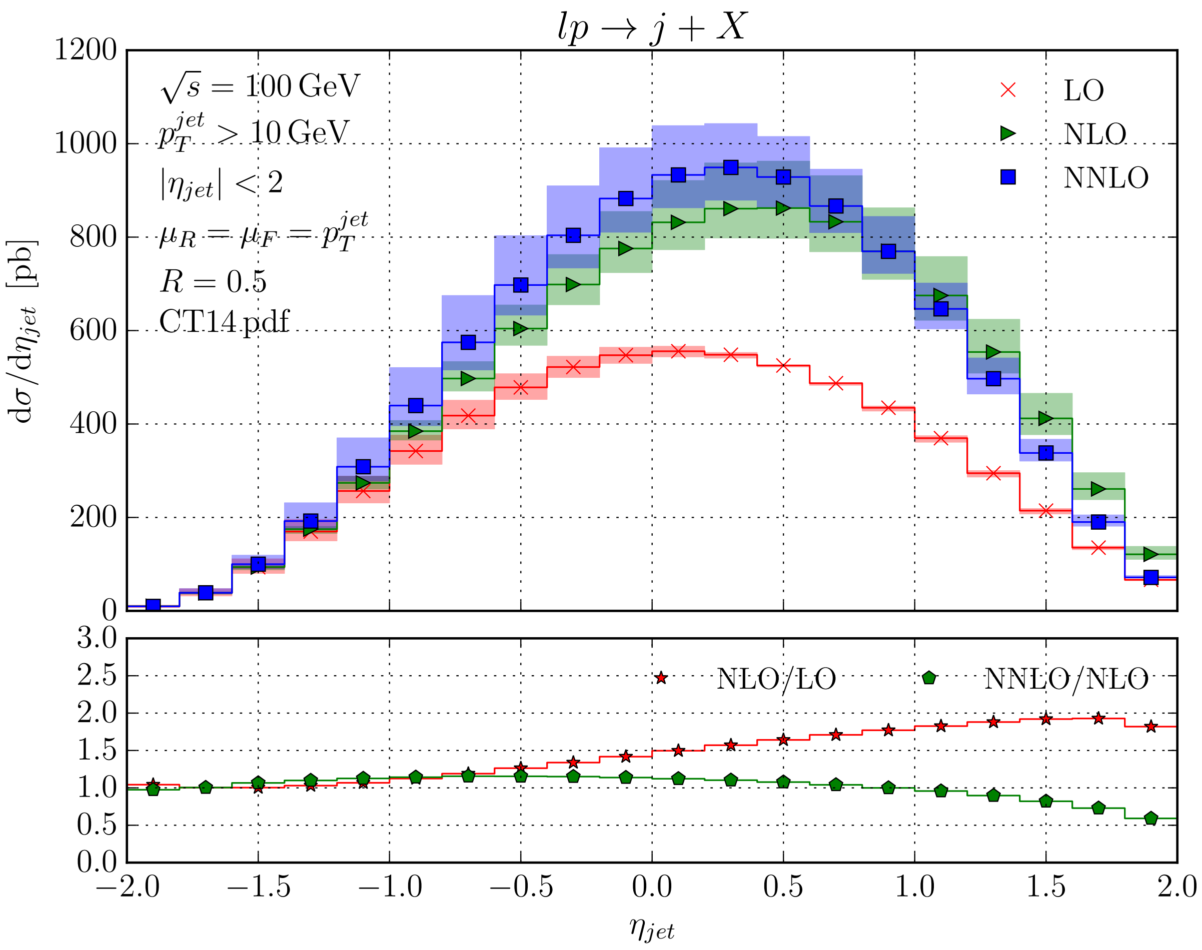

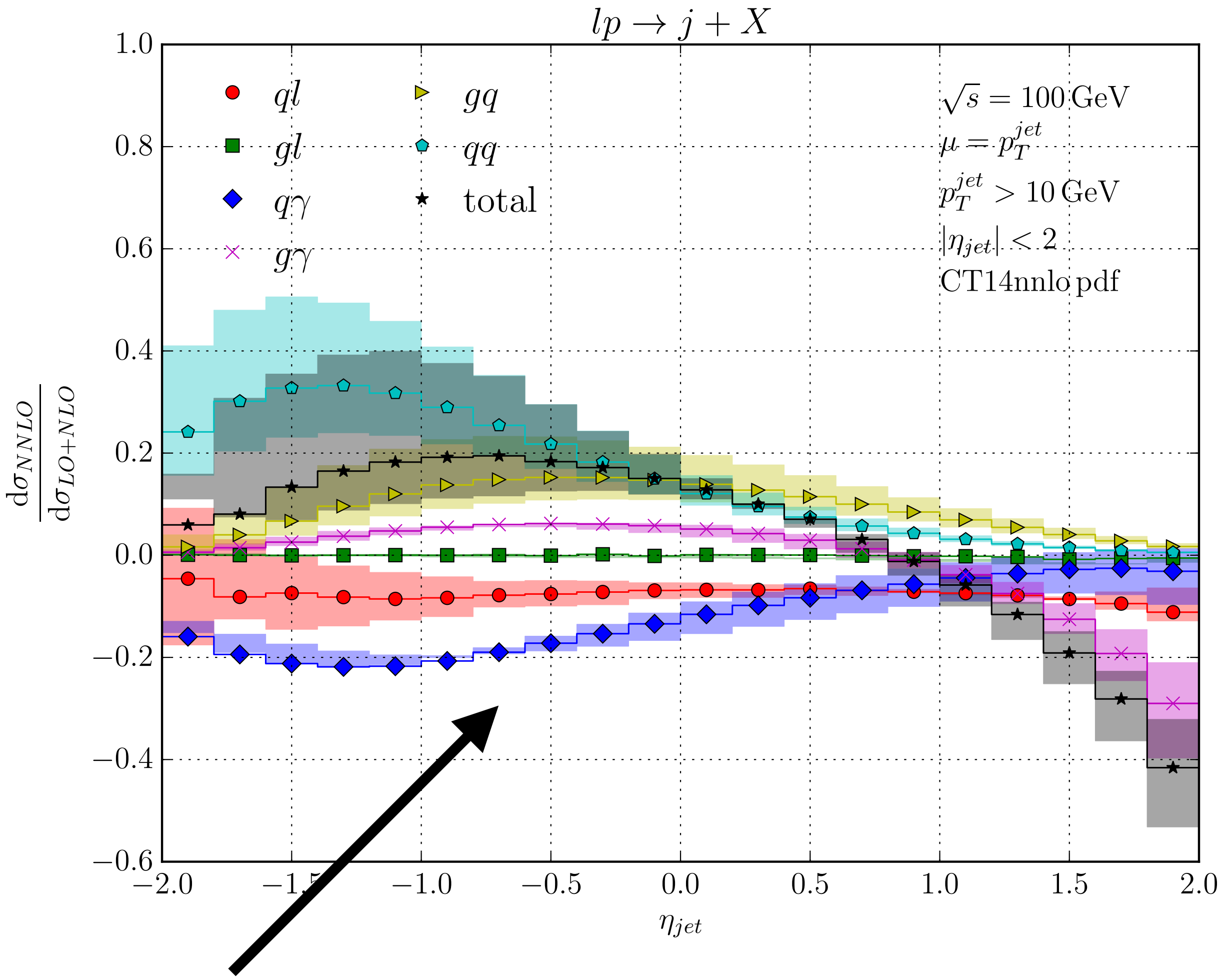

Different partonic channels dominate as a function of $\eta_{jet}$

Abelof, RB, Liu, Petriello 2016

# Scientific highlights

RB, Focke, Liu, Petriello PRL 2015

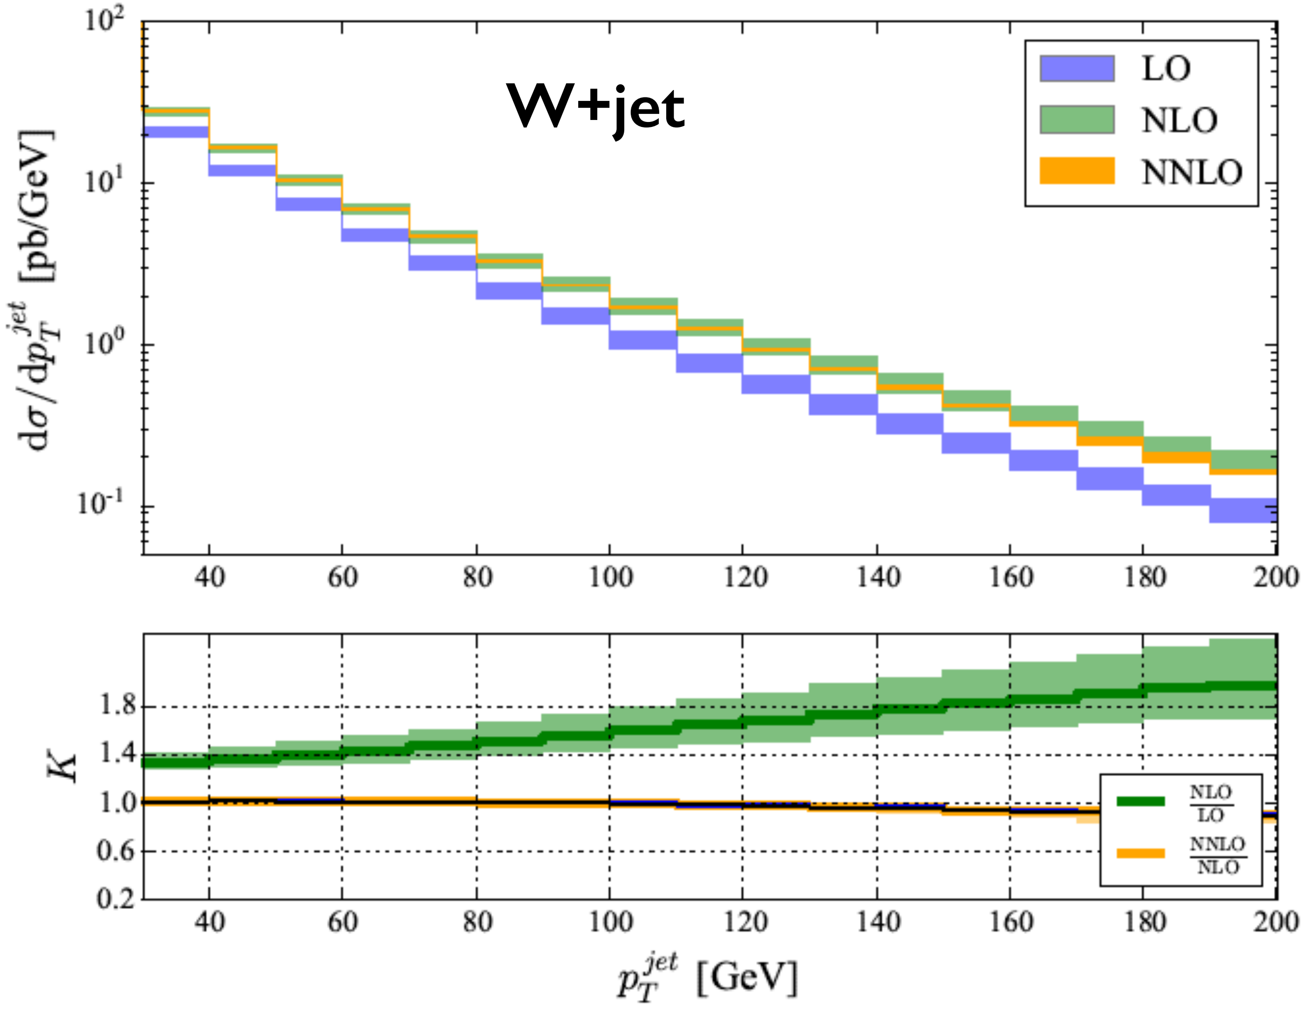

RB, Campbell, Ellis, Focke, Giele, Liu, Petriello PRL 2016

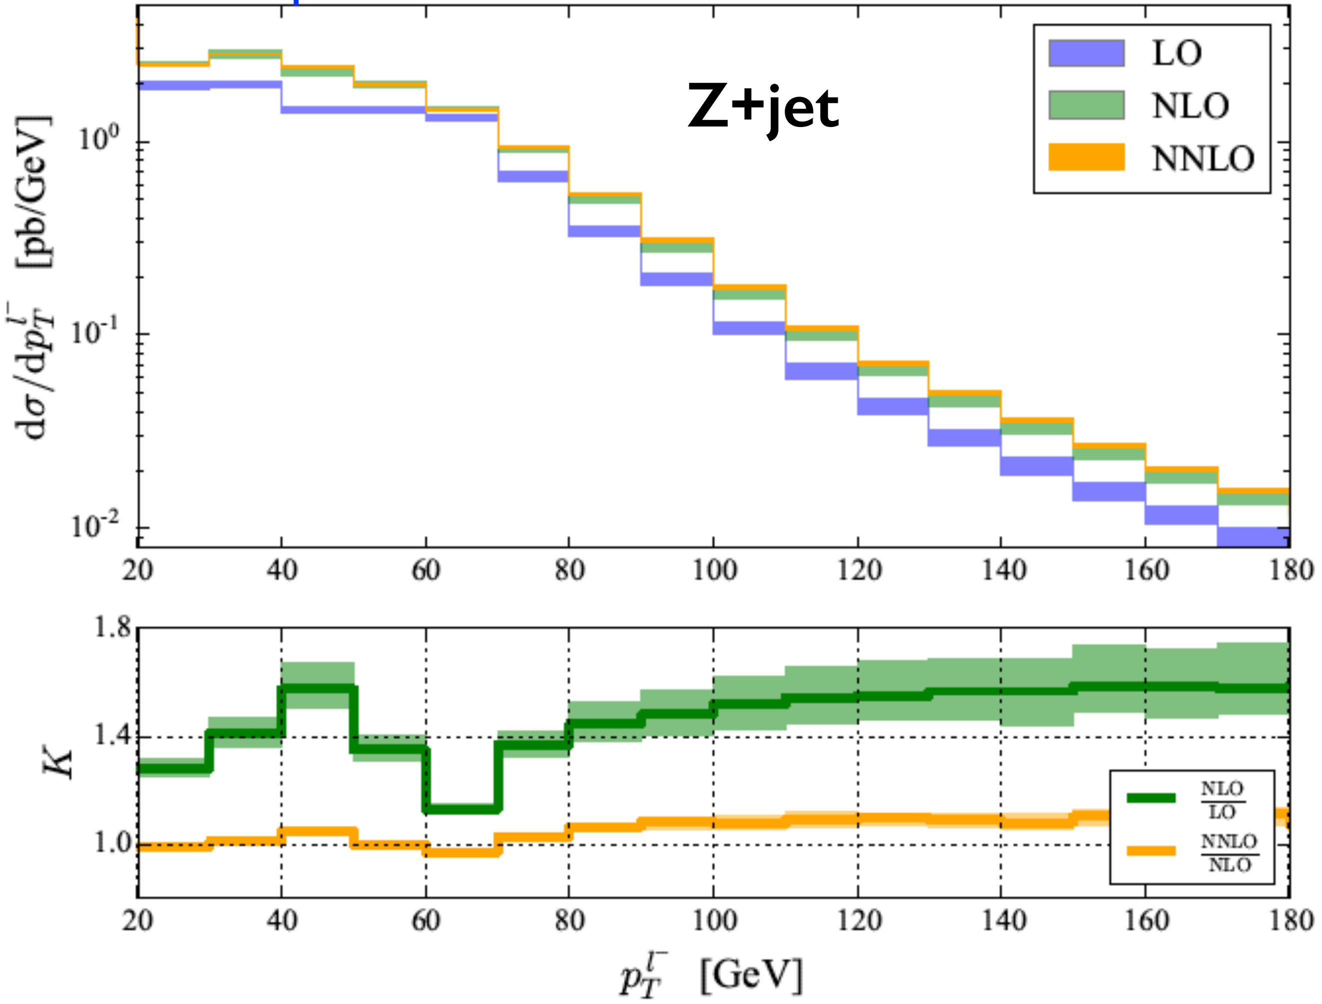

- First full calculations for W+jet@NNLO and Z+jet@NNLO; dramatic reduction of theory uncertainty upon inclusion of NNLO corrections

# Impact on the experimental program

- Both ATLAS and CMS have incorporated these calculations into their Run I and Run II analyses; a sampling of results is shown below

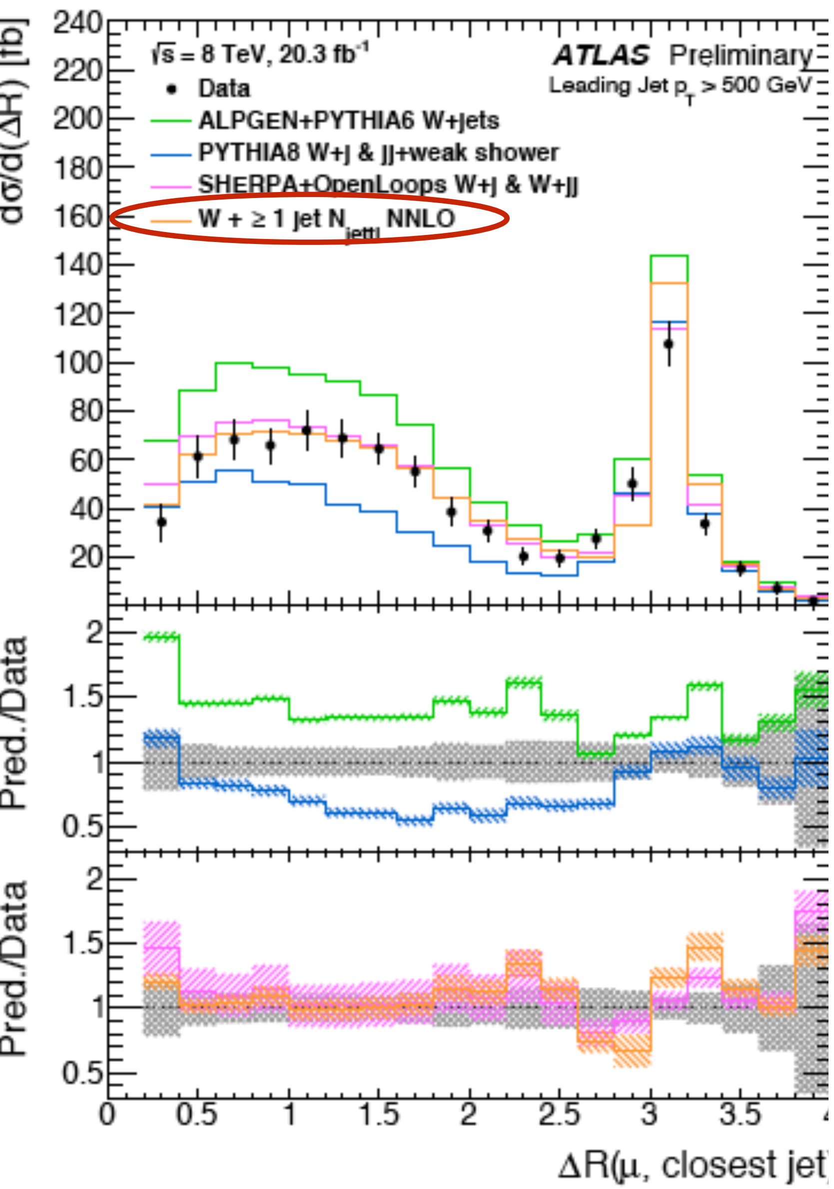

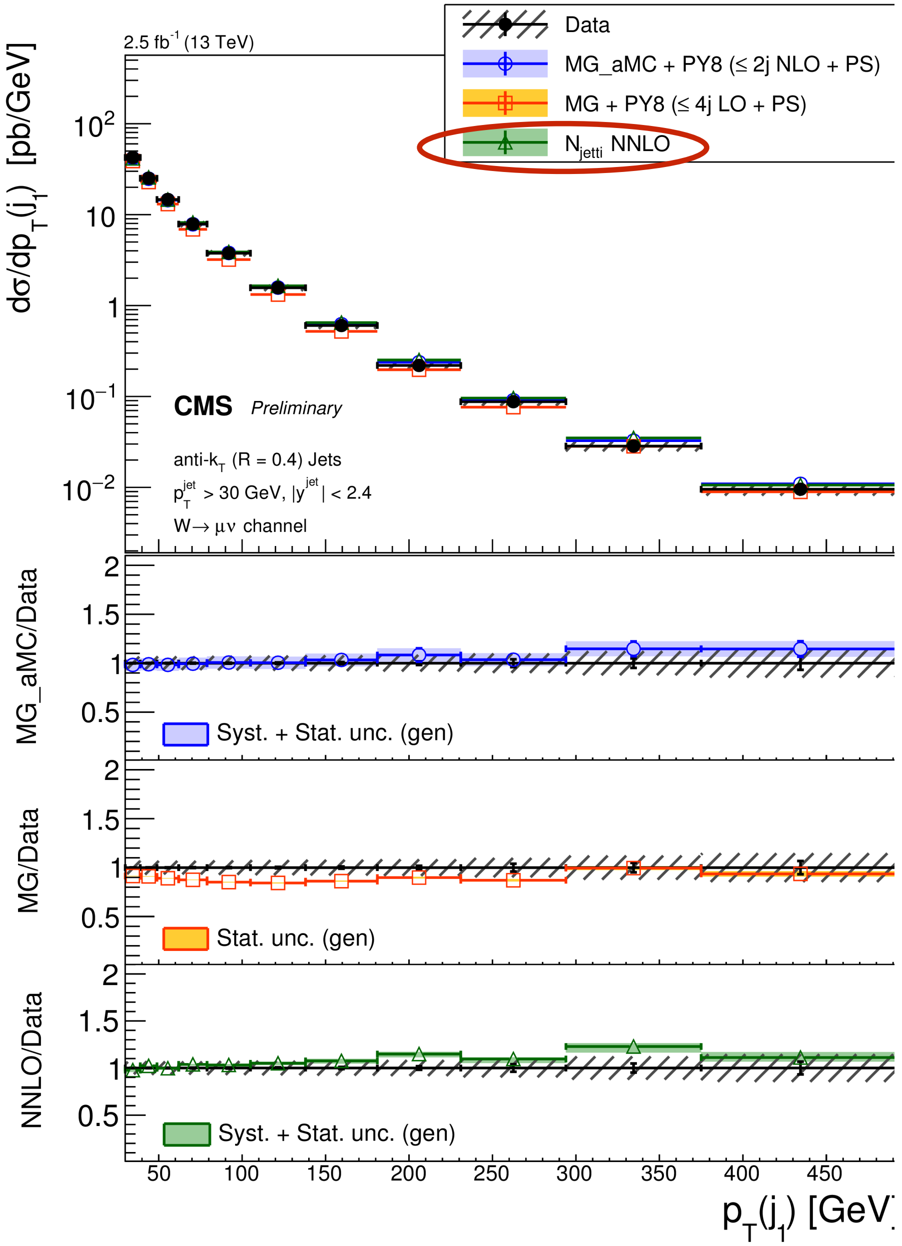

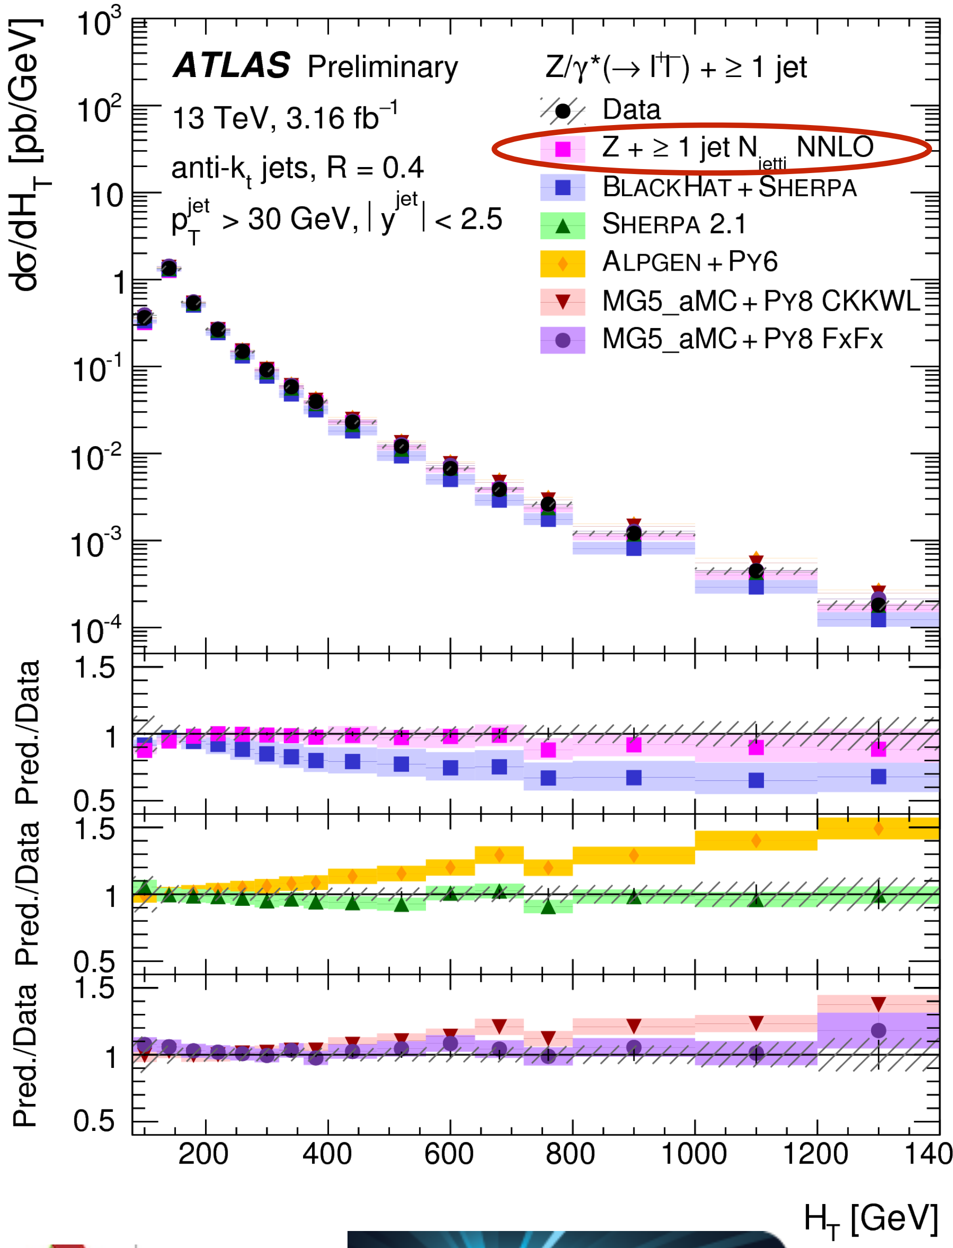

Excellent description of the entire suite of V+1-jet measurements!

# To-do-list

- Theoretical:
  - Theoretical refinements that will reduce $\tau_N^{cut}$ dependence
- Computational:
  - Everything compiled using GNU; can we use a different compiler?
  - Public release of V+jet, Higgs+jet codes
  - Public release of DISTRESS