Answers to Panel Questions

USQCD Collaboration: Norman Christ, Carleton DeTar, Will Detmold, Robert Edwards, Andreas Kronfeld, Christoph Lehner, Ethan Neil, Peter Petreczky

LQCD/USQCD Hardware Review
May 20–21, 2018
Brookhaven National Laboratory

Scientific Milestones: HEP

Scientific Milestones: Quark Flavor

- Quark flavor physics from 2013 WP & 2014 proposal:

Quantity	CKM element	(2013) Present expt. error	2007 forecast lattice error	(2013) Present lattice error	(forecast) 2014 lattice error	(forecast) 2018 lattice error
f_K/f_π	$\lvert V_{us}\rvert$	0.2%	0.5%	0.5%	0.3%	0.15%
$f_+^{K\pi}(0)$	$\lvert V_{us}\rvert$	0.2%	–	0.5%	0.35%	0.2%
f_D	$\lvert V_{cd}\rvert$	4.3%	5%	2%	1%	$< 1\%$
f_{D_s}	$\lvert V_{cs}\rvert$	2.1%	5%	2%	1%	$< 1\%$
$D \to \pi\ell\nu$	$\lvert V_{cd}\rvert$	2.6%	–	4.4%	3%	2%
$D \to K\ell\nu$	$\lvert V_{cs}\rvert$	1.1%	–	2.5%	2%	1%
$B \to D^*\ell\nu$	$\lvert V_{cb}\rvert$	1.3%	–	1.8%	1.5%	$< 1\%$
$B \to \pi\ell\nu$	$\lvert V_{ub}\rvert$	4.1%	–	8.7%	4%	2%
f_B	$\lvert V_{ub}\rvert$	9%	–	2.5%	1.5%	$< 1\%$
ξ	$\lvert V_{ts}/V_{td}\rvert$	0.4%	2-4%	4%	1.5%	$< 1\%$
ΔM_s	$\lvert V_{ts}V_{tb}\rvert^2$	0.24%	7–12%	11%	8%	5%
B_K	$\mathrm{Im}(V_{td}^2)$	0.5%	3.5–6%	1.3%	1%	$< 1\%$

- Progress in semileptonic *D* decays hampered by job security.

- Unexpected, but timely progress in rare semileptonic *B* decays.

Scientific Milestones: Quark Flavor

- Quark flavor physics from 2013 WP & 2014 proposal:

Quantity	CKM element	(2013) Present expt. error	2007 forecast lattice error	(2013) Present lattice error	(forecast) 2014 lattice error	(forecast) 2018 lattice error	actual 2018 lattice error	when achieved
f_K/f_π	$\lvert V_{us}\rvert$	0.2%	0.5%	0.5%	0.3%	0.15%	0.15%	end 2017
$f_+^{K\pi}(0)$	$\lvert V_{us}\rvert$	0.2%	–	0.5%	0.35%	0.2%	$0.27 \to 0.19\%$	2013, 2015, 2018
f_D	$\lvert V_{cd}\rvert$	4.3%	5%	2%	1%	$< 1\%$	0.25%	end 2017
f_{D_s}	$\lvert V_{cs}\rvert$	2.1%	5%	2%	1%	$< 1\%$	0.16%	end 2017
$D \to \pi\ell\nu$	$\lvert V_{cd}\rvert$	2.6%	–	4.4%	3%	2%	4.4%	2011
$D \to K\ell\nu$	$\lvert V_{cs}\rvert$	1.1%	–	2.5%	2%	1%	2.5%	2010
$B \to D^*\ell\nu$	$\lvert V_{cb}\rvert$	1.3%	–	1.8%	1.5%	$< 1\%$	1.4%	2014
$B \to \pi\ell\nu$	$\lvert V_{ub}\rvert$	4.1%	–	8.7%	4%	2%	4.3%	2015 expt$\oplus$QCD
f_B	$\lvert V_{ub}\rvert$	9%	–	2.5%	1.5%	$< 1\%$	0.74%	end 2017
ξ	$\lvert V_{ts}/V_{td}\rvert$	0.4%	2-4%	4%	1.5%	$< 1\%$	1.5%	2016
ΔM_s	$\lvert V_{ts}V_{tb}\rvert^2$	0.24%	7–12%	11%	8%	5%	8.3%	2016
B_K	$\mathrm{Im}(V_{td}^2)$	0.5%	3.5–6%	1.3%	1%	$< 1\%$	1.37%	2011, 2014, 2015

- Progress in semileptonic *D* decays hampered by job security.

- Unexpected, but timely progress in rare semileptonic *B* decays.

Scientific Milestones: Muon g–2

- In 2013, it was unclear how quickly the calculations for the two hadronic contributions would develop, so no formal forecasts were formulated.
- Hadronic vacuum polarization since 2013:
  - Standard milestone of first calculation with full error budget has been passed, by more than one effort (from USQCD, with more worldwide).
  - Replacing conflicting e^+e^- data with lattice QCD underway.
  - Lattice-QCD only calculation with E989 uncertainties a few years away.
- Hadronic light-by-light since 2017:
  - Rapid development starting last year [arXiv:1705.01067] makes the first calculation with a full error budget likely in **2019**.

Scientific Milestones: BSM
(based on 2016 presentation)

Detailed scientific targets: composite Higgs

1. Identify the most likely candidates in the spectrum for LHC discovery in the theories we are currently studying. Compute matrix elements for their decay widths, and study the most promising search channels (e.g. diphoton for spin-0 resonances.)

2. Study the emergent low-energy effective theory of pions and the light 0++ "sigma". Calculate the interactions between these states, and see if they match predictions from known EFTs such as the linear sigma model or chiral perturbation theory.

3. Calculate anomalous dimensions and "baryon"-to-vacuum matrix elements, for theories of partial compositeness. Begin with theories that have proposed UV completions: SU(3) with 4 light fermions, and SU(4) with fermions in fundamental and antisymmetric reps.

(new prelim. results - see BSM overview slides)

Detailed scientific targets: composite Higgs

1. Identify the <u>most likely candidates in the spectrum for LHC discovery</u> in the theories we are currently studying. Compute matrix elements for their decay widths, and study the most promising search channels (e.g. diphoton for spin-0 resonances.)

arXiv:1710.00806*

arXiv:1801.05809

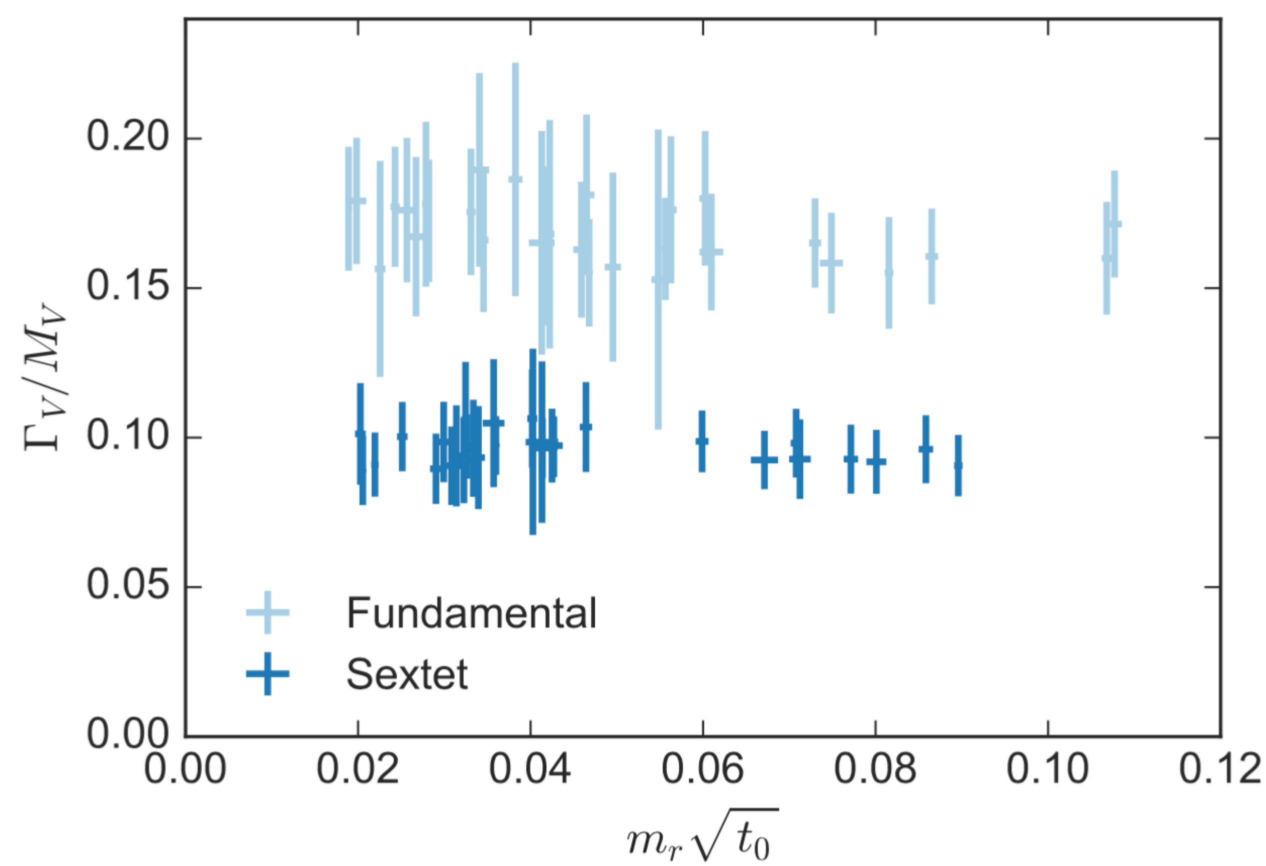

*Mass-to-width ratio for SU(4) composite Higgs model from lattice + KSRF sum rules. Predict two new vector resonances with characteristic ratio of widths (somewhat narrower than QCD rho.)

(new prelim. results - see BSM overview slides)

Detailed scientific targets: composite dark matter

1. Calculate finite-temperature transition properties (transition order, latent heat, etc.) in a candidate dark matter theory, e.g. SU(4) “stealth DM”. Match on to gravitational wave predictions, predict possible signals for future gravitational wave observatories.

2. Determine the binding energy of light nuclei in SU(4) gauge theory, extending known results for QCD and for SU(2). Use the results to study whether large-N expansion is effective here.

3. Determine meson electromagnetic form factors in SU(4) gauge theory. Test the usefulness of vector-meson dominance in theories other than QCD. Make predictions for collider production rates and for indirect detection via dark matter annihilation.

Detailed scientific targets: composite dark matter

1. Calculate finite-temperature transition properties (transition order, latent heat, etc.) in a candidate dark matter theory, e.g. SU(4) “stealth DM”. Match on to gravitational wave predictions, predict possible signals for future gravitational wave observatories.

arXiv:1802.09644*

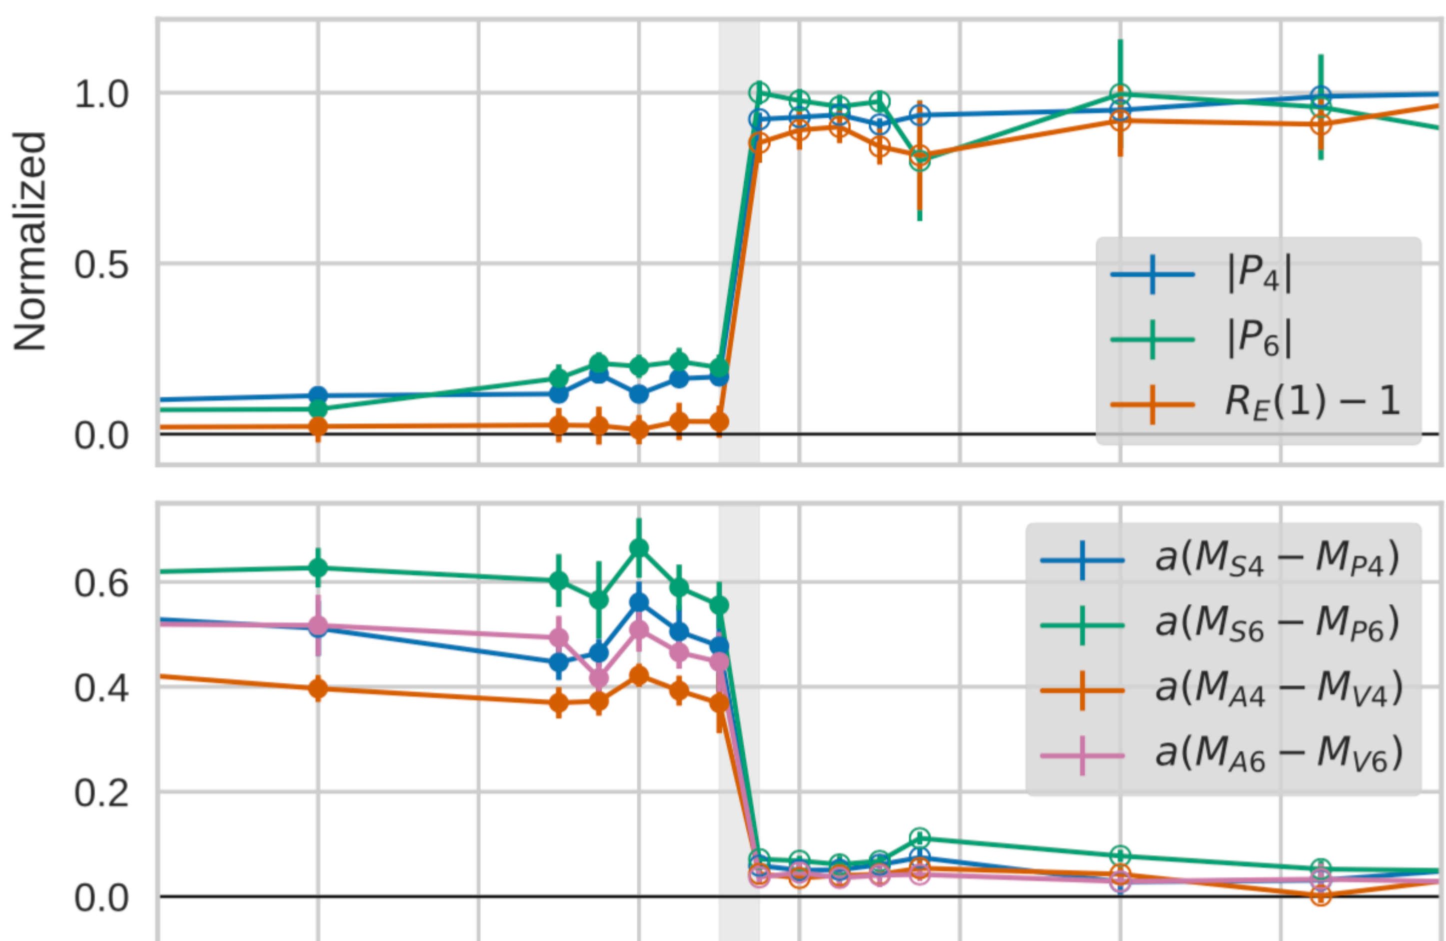

*Study of phase transition in SU(4) two-representation model; found to be strongly first-order. Follow-up study of transition properties for GW signal in future.

Detailed scientific targets: lattice SUSY

1. Compute the conformal dimensions of various primary operators in superconformal phase of N=4 Yang-Mills, e.g. the Konishi - lightest flavor singlet scalar operator. Compare non-perturbative results to conformal bootstrap bounds, and to large-N prediction (Bethe ansatz).

2. Investigate the lattice beta function for N=4 Yang-Mills using Monte Carlo Renormalization Group methods. Determine what fine tuning is needed to restore full supersymmetry in continuum limit.

3. Study the scaling of the static potential at large N; test for evidence of Maldacena scaling.

4. Compute the W boson and monopole masses on the Coulomb branch of the theory. In combination with (1), provide a direct test of the weak-strong (S) duality conjecture. Also, test electric/magnetic line operator duality.

5. Make contact with string/supergravity theory through holographic connections. Pursue any lattice results at finite N where corrections to leading SUGRA results from the string-theory side can be studied.

Detailed scientific targets: lattice SUSY

arXiv:1709.07025

arXiv:1710.06398*

*Two-dimensional maximally supersymmetric Yang Mills, non-perturbative lattice results at large N confirm leading supergravity prediction. See BSM overview slides for preliminary d=3 result.

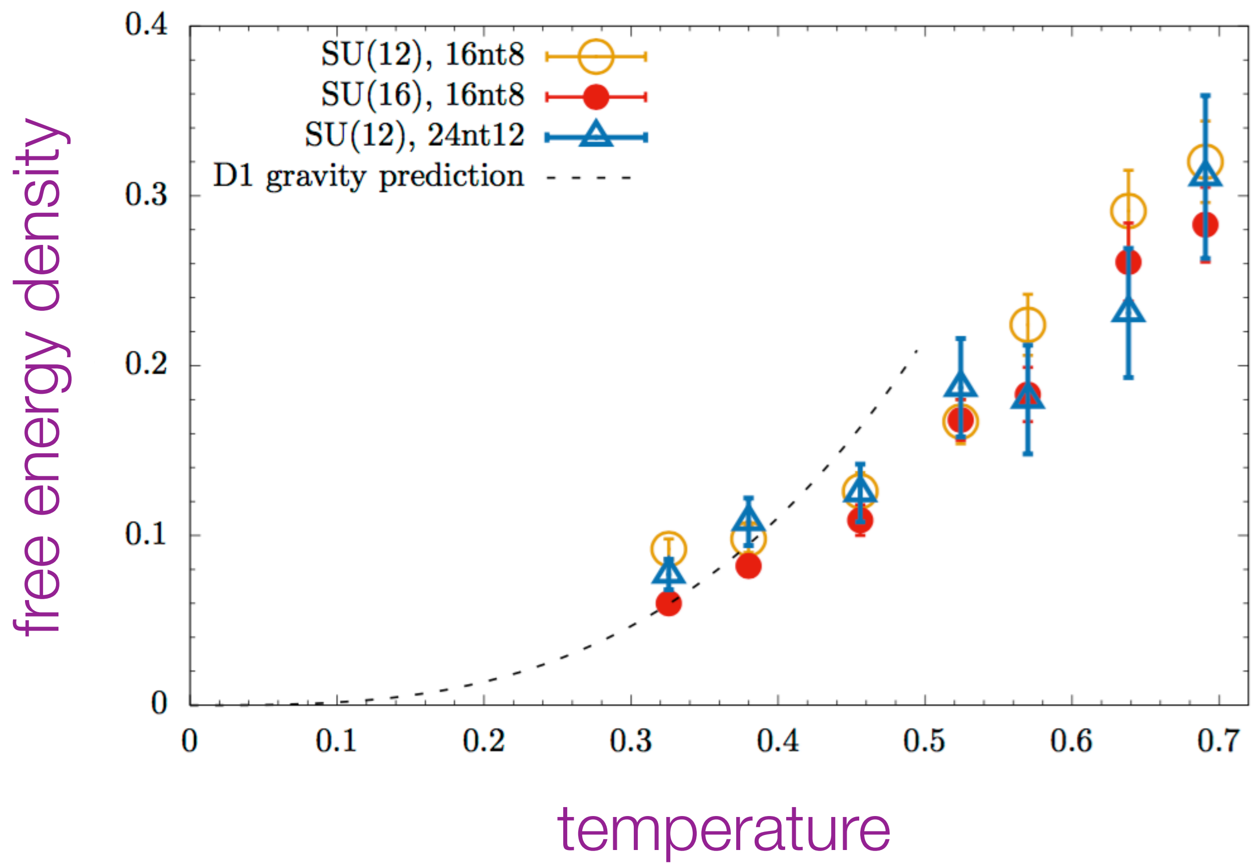

Scientific Milestones: Cold NP (compare 2016 & 2018)

Precision Goals Table

2016 edition

- Quantities with currently quantifiable uncertainties and goals

Quantity	Current uncertainty	Uncertainty Goal	Impact/Target
g_A	5%	3%->1%	Benchmark of LQCD; Neutrino-nucleus X-secs V_{ud} given high enough precision
L_u, L_d	~20%	5%	Understanding the spin of the proton
g_S, g_T	5%,20%	10%	Ultracold neutron experimental searches for BSM interactions in neutron decay
$a_{\pi\pi}^{(I=2)}$	1%	✓	More precise than experiment/phenomenology
$\langle N \vert \bar{s}s \vert N \rangle$	25%	10%	Input for dark matter direct detection experiments; mu2e conversion
$\langle x \rangle$	~15%	3%	Aim for ab initio input to PDFs (USQCD goal)
$\langle r^2 \rangle_p$	~25%	2%	Impact proton radius puzzle

Precision Goals and Progress

2018 edition (updated)

Progress highlighted in red

Quantity	Current uncertainty	Uncertainty Goal	Impact/Target
g_A	**3%**	<1%	Benchmark of LQCD; Neutrino-nucleus X-secs V_{ud} given high enough precision
L_u, L_d	~20%,~10%	5%	Understanding the spin of the proton
g_S, g_T	**12%**, 5%	10%, 3%	Ultracold neutron experimental searches for BSM interactions in neutron decay, DM
$a_{\pi\pi}^{(I=2)}$	1%	✓	More precise than experiment/phenomenology
$\langle N\|\bar{s}s\|N\rangle$	25%	10%	Input for dark matter direct detection experiments; mu2e conversion
$\langle x\rangle$	~15%	5%	Aim for ab initio input to PDFs (USQCD goal)
$\langle r^2\rangle_p$	~10%	2%	Impact proton radius puzzle
L_1, L_{1A}	**20%**	5%	pp fusion; next generation neutrino detectors

PDF goals

- Community white paper
  [Lin et al, Prog. Part. Nucl. Phys. 100 (2018) 107-160]
- Fake lattice data experiments
  - Given various improvements in LQCD calculate, what is the impact

Scenario	$\delta_L^{(i)}$ for unpolarized moments				
	$\langle x\rangle_{u^+}$	$\langle x\rangle_{d^+}$	$\langle x\rangle_{s^+}$	$\langle x\rangle_g$	$\langle x\rangle_{u^+-d^+}$
Current	$\sim 16\%$	$\sim 30\%$	$\sim 45\%$	$\sim 13\%$	$\sim 60\%$
A	3%	3%	5%	3%	5%
B	2%	2%	4%	2%	4%
C	1%	1%	3%	1%	3%

  - Impact on large x PDF
  - Impact on polarised gluons

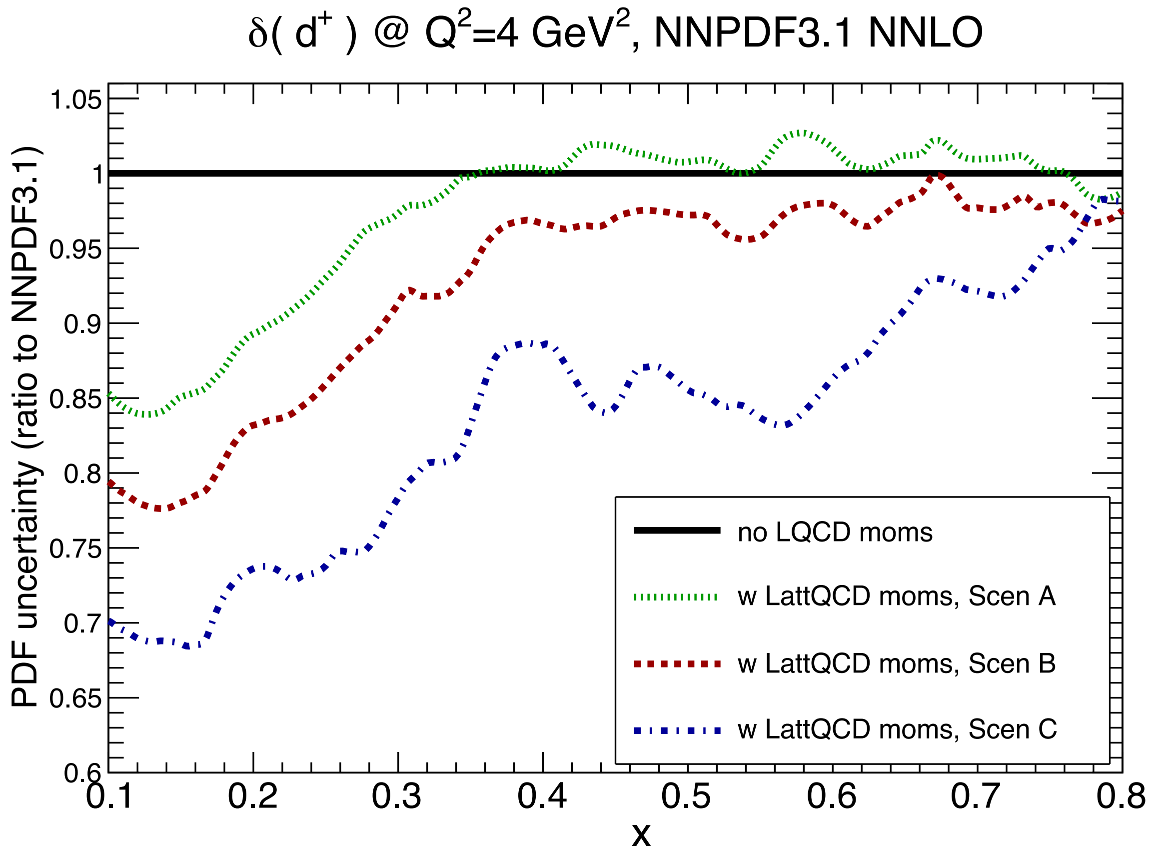

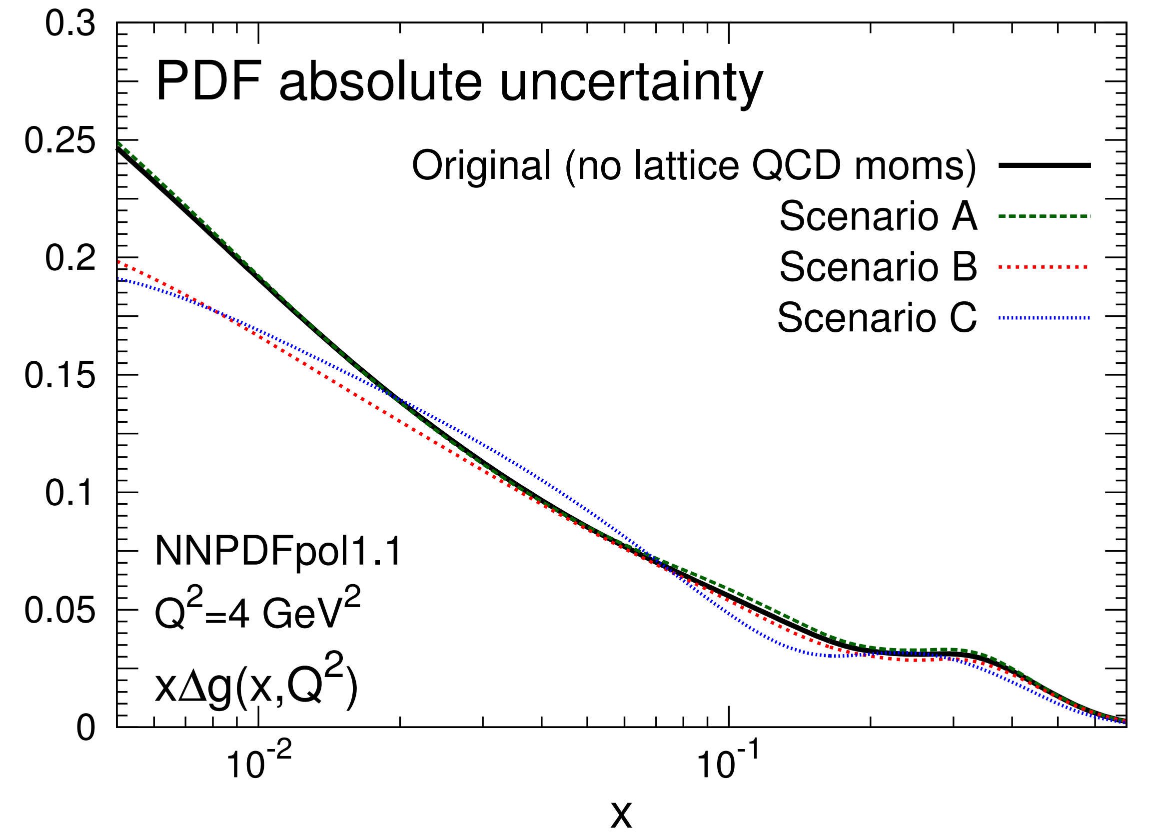

FLAG 2016
Flavour Lattice Averaging Group

Eur. Phys. J. C (2017) 77:112
DOI 10.1140/epjc/s10052-016-4509-7

THE EUROPEAN PHYSICAL JOURNAL C

CrossMark

Review

Review of lattice results concerning low-energy particle physics

Flavour Lattice Averaging Group (FLAG)

S. Aoki[1], **Y. Aoki**[2,3,17], **D. Bečirević**[4], **C. Bernard**[5], **T. Blum**[3,6], **G. Colangelo**[7], **M. Della Morte**[8,9], **P. Dimopoulos**[10,11], **S. Dürr**[12,13], **H. Fukaya**[14], **M. Golterman**[15], **Steven Gottlieb**[16], **S. Hashimoto**[17,18], **U. M. Heller**[19], **R. Horsley**[20], **A. Jüttner**[21,a], **T. Kaneko**[17,18], **L. Lellouch**[22], **H. Leutwyler**[7], **C.-J. D. Lin**[22,23], **V. Lubicz**[24,25], **E. Lunghi**[16], **R. Mawhinney**[26], **T. Onogi**[14], **C. Pena**[27], **C. T. Sachrajda**[21], **S. R. Sharpe**[28], **S. Simula**[25], **R. Sommer**[29], **A. Vladikas**[30], **U. Wenger**[7], **H. Wittig**[31]

- Next edition of FLAG (Flavor Lattice Averaging Group) review will contain various nucleon quantities
  - Axial Charge
  - Scalar and Tensor “Charges”
- Develop robust community consensus on calculations

Scalar and tensor charges

Thanks to Huey-Wen Lin

- Isovector scalar and tensor charges @ scale 2 GeV

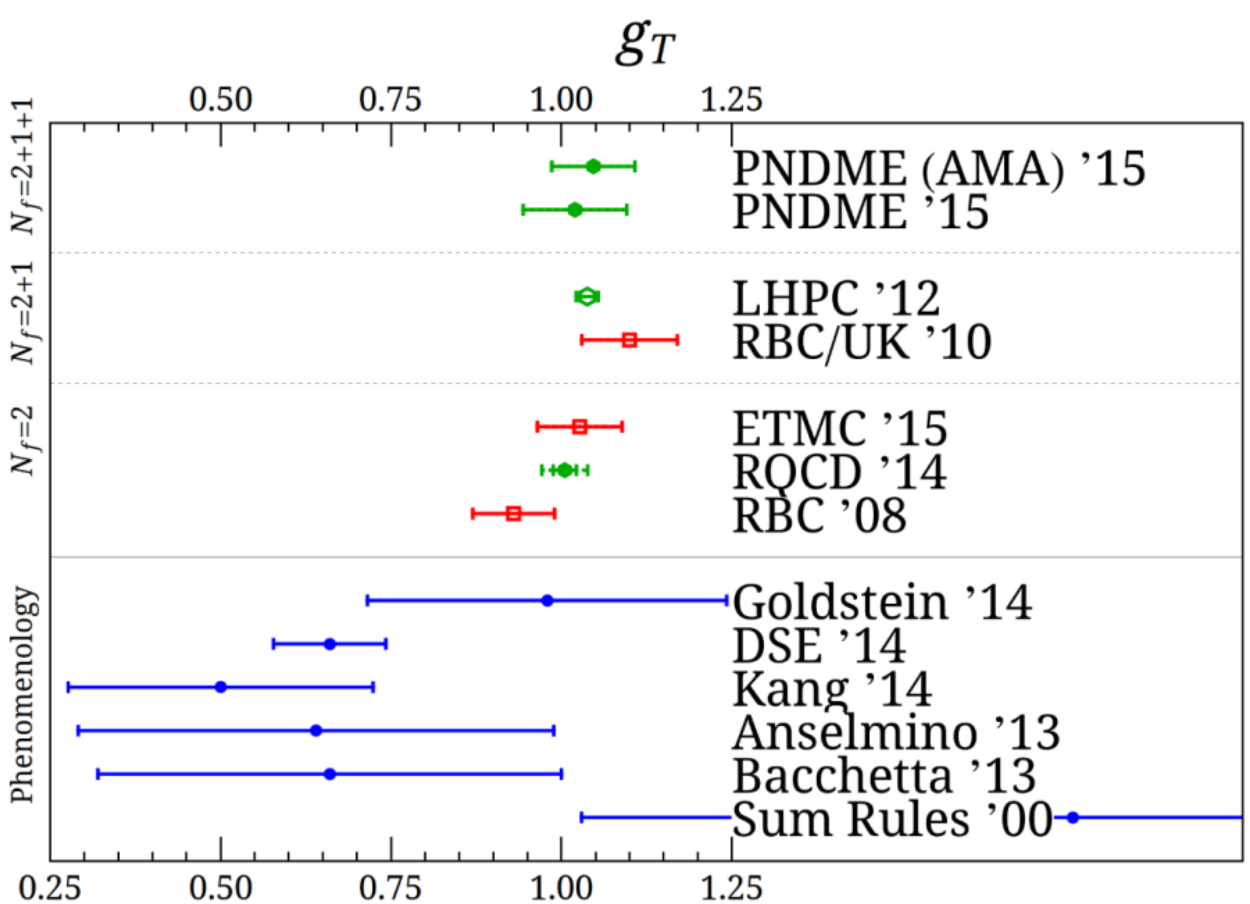

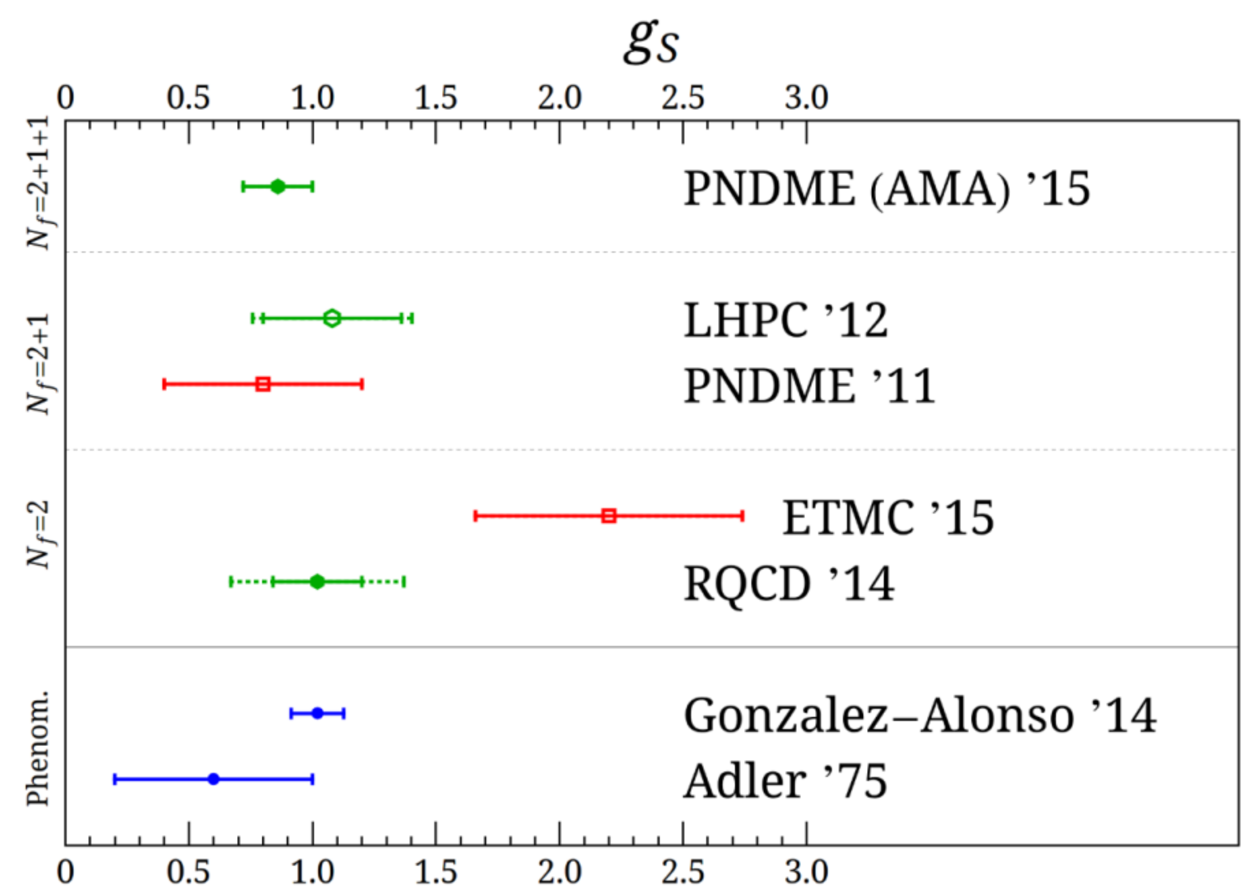

Collaboration	Ref.	publication status	N_f	chiral extrapolation	continuum extrapolation	finite volume	excited state	renormalization	g_T
PNDME'15	This work	P	2+1+1	★	★	★	★	★	1.020(76)[a]
ETMC'13	[30]	C	2+1+1	■	○	○	■	★	1.11(3)[b]
LHPC'12	[28]	A	2+1	★	○	★	○	★	1.037(20)[c]
RBC/UKQCD'10	[29]	A	2+1	○	■	★	★	★	1.10(7)[d]
RQCD'14	[31]	P	2	★	★	★	○	★	1.005(17)(29))[e]
ETMC'13	[30]	C	2	★	■	○	■	○	1.114(46)[f]
RBC'08	[32]	P	2	■	■	★	■	★	0.93(6)[g]

Scalar and tensor charges

Thanks to Huey-Wen Lin

- Isovector scalar and tensor charges @ scale 2 GeV

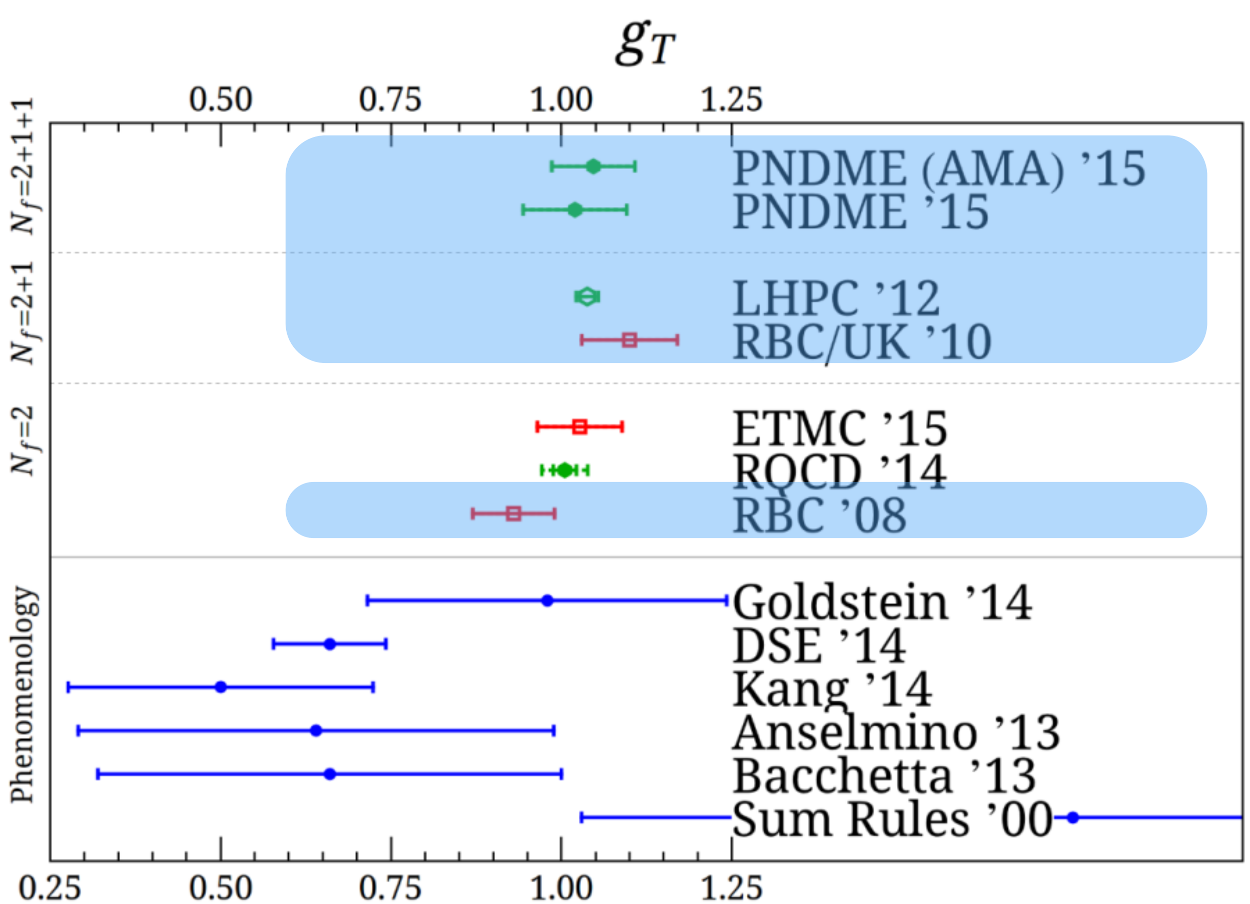

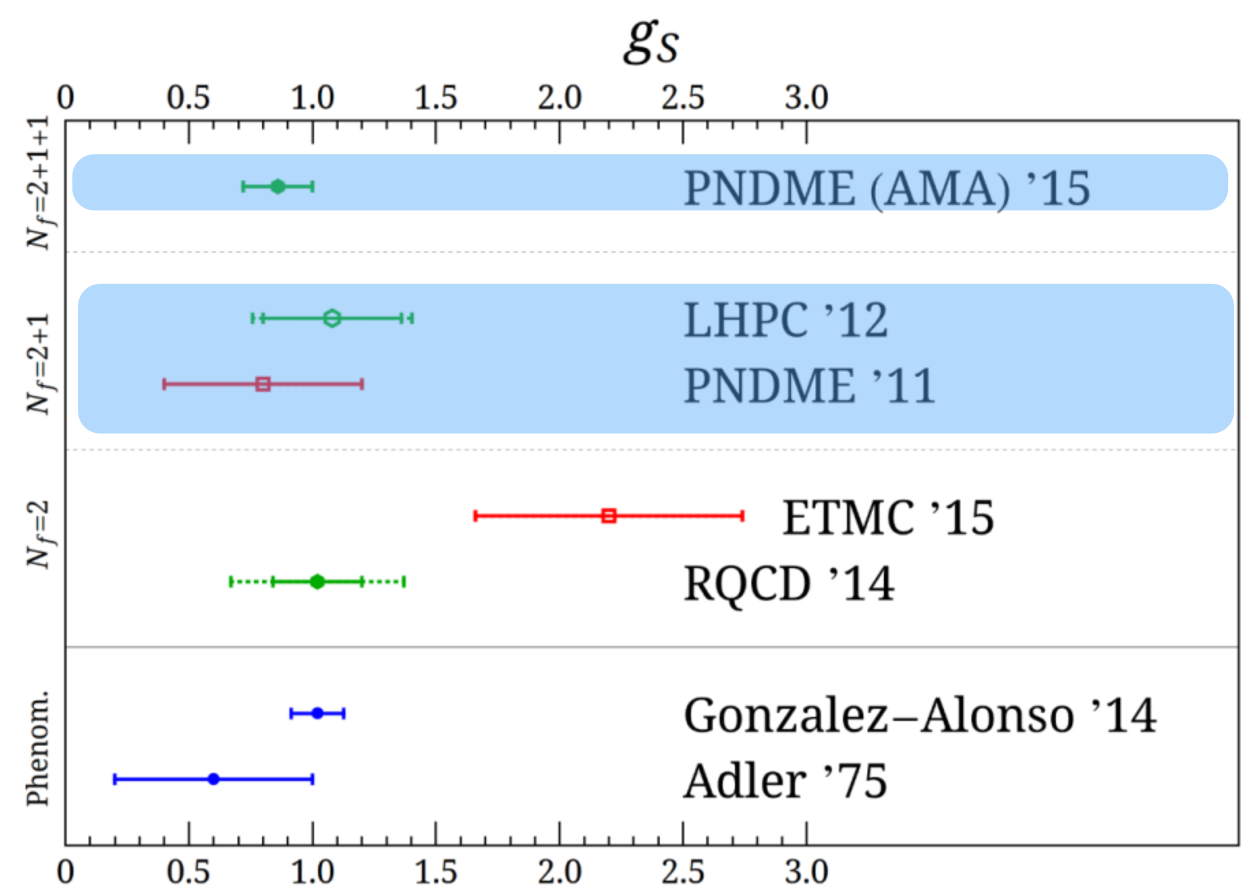

Collaboration	Ref.	publication status	N_f	chiral extrapolation	continuum extrapolation	finite volume	excited state	renormalization	g_T
PNDME'15	This work	P	2+1+1	★	★	★	★	★	1.020(76)[a]
ETMC'13	[30]	C	2+1+1	■	○	○	■	★	1.11(3)[b]
LHPC'12	[28]	A	2+1	★	○	★	○	★	1.037(20)[c]
RBC/UKQCD'10	[29]	A	2+1	○	■	★	★	★	1.10(7)[d]
RQCD'14	[31]	P	2	★	★	★	○	★	1.005(17)(29))[e]
ETMC'13	[30]	C	2	★	■	○	■	○	1.114(46)[f]
RBC'08	[32]	P	2	■	■	★	■	★	0.93(6)[g]

Sigma term and strangeness content

Thanks to Phiala Shanahan

- Nucleon sigma term and strangeness content of nucleon

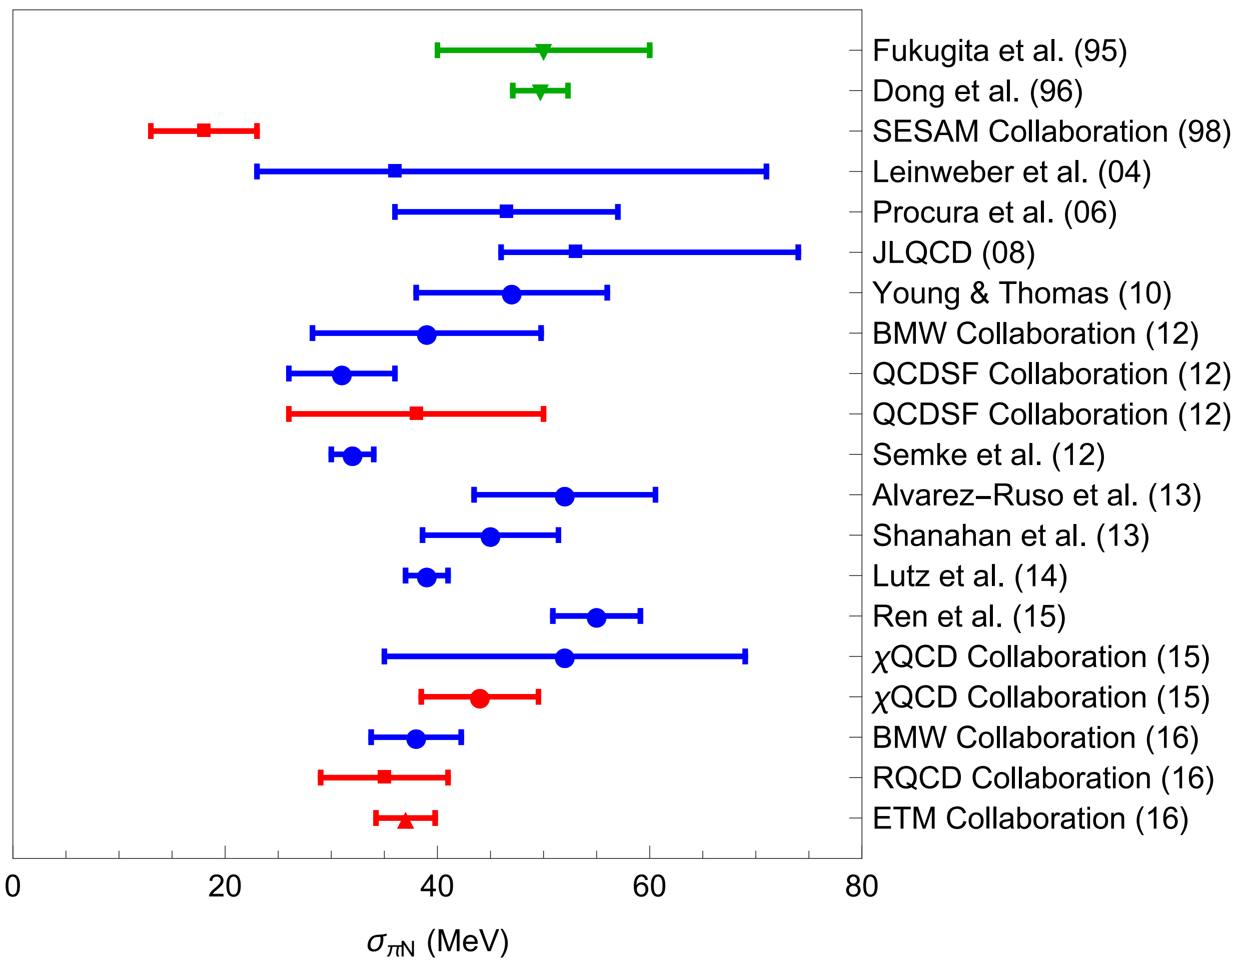

(a) $\sigma_{\pi N} = \frac{m_u + m_d}{2}\langle N|\bar{u}u + \bar{d}d|N\rangle$

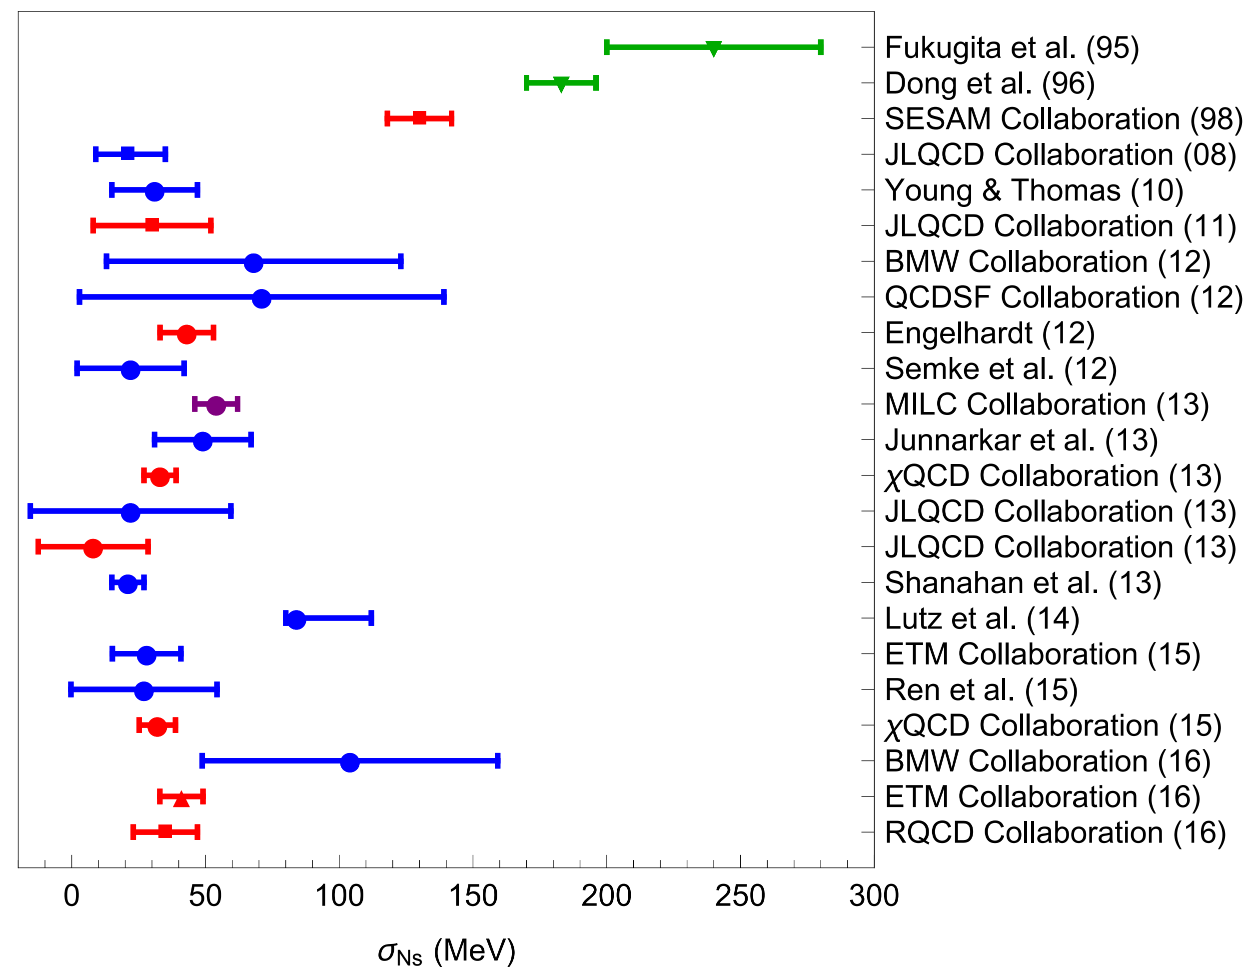

(b) $f_s^N = \frac{\sigma_{sN}}{m_N}$

Sigma term and strangeness content

Thanks to Phiala Shanahan

- Nucleon sigma term and strangeness content of nucleon

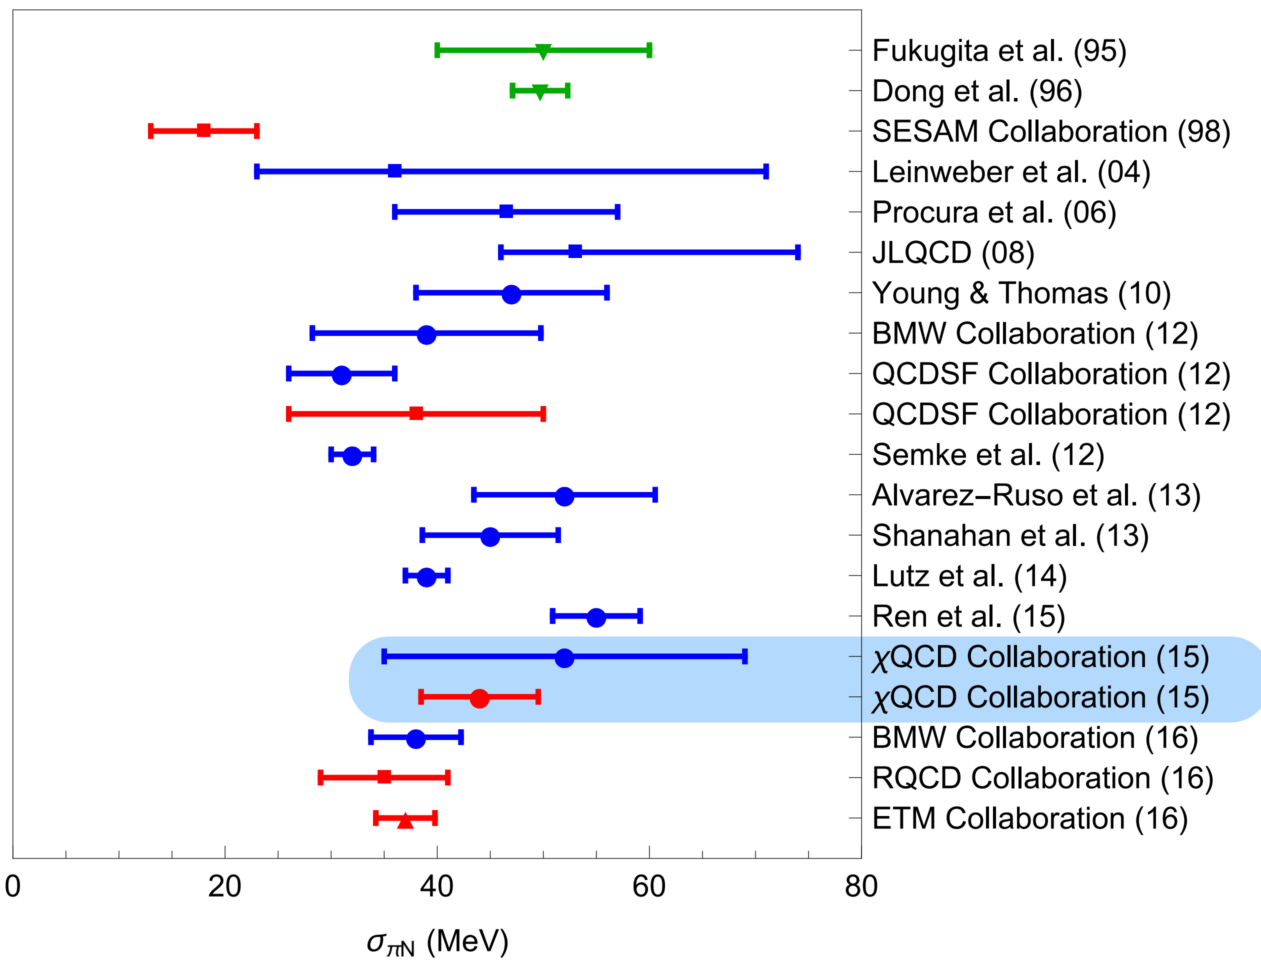

(a) $\sigma_{\pi N} = \frac{m_u+m_d}{2}\langle N|\bar{u}u + \bar{d}d|N\rangle$

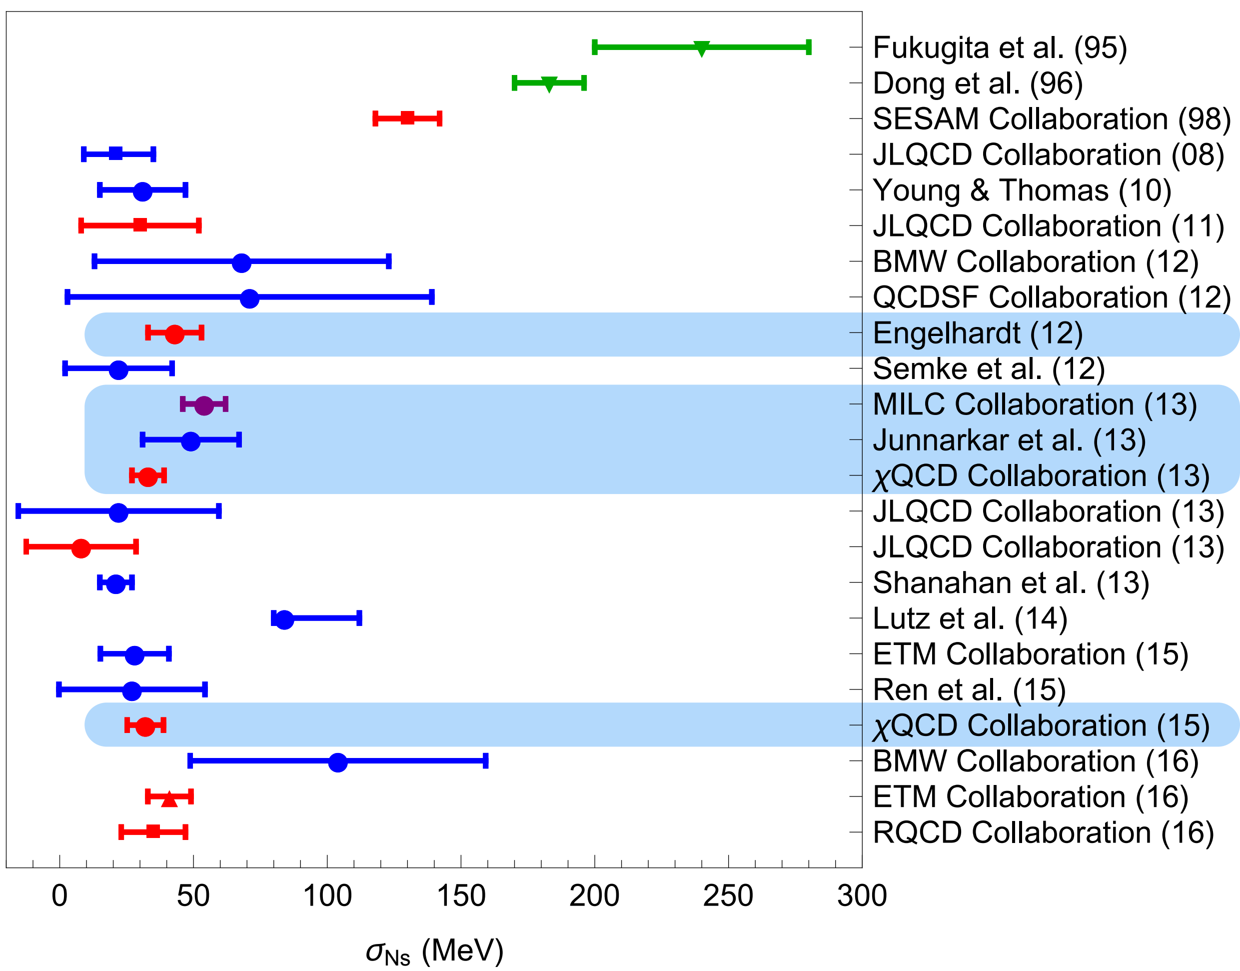

(b) $f_s^N = \frac{\sigma_{sN}}{m_N}$

Scientific Milestones: Thermodynamics
(*cf*. 2016 Review, Swagato Mukherjee)

USQCD hot-dense LQCD goals (~2020)

provide urgently needed QCD inputs & baselines
essential for dynamical modeling of BES energies

bulk thermodynamics @ $\mu_B > 0$

- equation of state
- QCD phase boundary
- equilibrium baseline for higher cumulants
- freeze-out conditions
- ...

integral part of

BEST COLLABORATION

Topical Collaboration
in nuclear theory
2016-2020
DOE, NP

USQCD hot-dense LQCD goals (~2020)

provide urgently needed QCD inputs & baselines essential for dynamical modeling of BES energies

bulk thermodynamics @ $\mu_B>0$

- equation of state
- QCD phase boundary
- equilibrium baseline for higher cumulants
- freeze-out conditions
- ...

2020 goal needs 8th order in μ_B; Peter showed results for 4th order; feasibility depends on, e.g., ALCC.

Other main goal for “now” was T dependence of quarkonium spectral function: presented last week at Quark Matter ’18; arXiv soon.

Pros and Cons of Institutional Clusters

Pros to IC (science perspective)

- Flexibility: that (not surprisingly) USQCD-software supports/pursues
  - BNL worked closely with LQCD extension II - good outcome for USQCD.
  - BNL has 3 architectures, GPUs, KNLs, Skylakes (coming).
  - BNL provisioned for our needs (e.g., sufficient network support for expected jobs sizes, sufficient I/O.
- Large facility can absorb/support fluctuations in USQCD usage demand.
- (Comparatively) large support staff - in-depth knowledge (disk systems, network, tape, etc.)

Cons to ICs (science perspective)

- No guarantee that all IC procurements will be as pleasant as 2017 BNL-USQCD experience:
  - LQCD may be too small to ensure adequate customization.
- LQCD is adept at migrating to new/faster/better systems. Will lab-wide IC boundary conditions keep us so nimble?
  - (This is the argument for why dedicated systems are optimal for a large-enough user base.)
- Allocations & account-year not matched to USQCD.
- [Obviously, many of these pro/cons not unique to USQCD/LQCD experiences]

NP-funded SciDAC-4

ASCR/NP SciDAC-4

Computing the Properties of Matter with Leadership Computing Resources

Co-PIs: Will Detmold (MIT), Balint Joo (JLab), Swagato Mukherjee (BNL)

Senior Investigators:

Andrei Alexandrou (GWU)
Saman Amarasinghe (MIT)
Alexei Bazavov (MSU)
Kate Clark (NVIDIA)
Rob Fowler (UNC)
Dhiraj Kalamkar (Intel)
Xu Liu (W&M Computer Sci)
Kostas Orginos (W&M Phys)
Sergey Panitkin (BNL)
Andrew Pochinsky (MIT)
Martin Savage (UW)
Frank Winter (JLab)
Boram Yoon (LANL)

Budget: FY2017 - 2022: $4.25M (ASCR) + $4M (NP)

Main areas of activity

- Gauge generation (co-PI: Balint Joo)
  - Clover on emerging LCF-s (Summit, NERSC-9 - and originally Aurora), GPU/KNL inverters
  - Frank Winter: JIT
  - Boram Yoon - HMC integrators
  - Subcontracts:
    - Kostas Orginos - 1 flavor methods for Clover, determinant reweighing methods
    - Xu Liu - memory optimizations (QPerf and HPC Toolkit)
    - Rob Fowler- QUARC/DSL interface to Clang - automatic code generation

- Correlation functions/contractions (co-PI: Will Detmold)
  - Saman Amarasinghe - (TACO) code generation, auto-tuning for contractions (& gauge gen)
  - Andrew Pochinsky - Halide for QCD
  - Kenneth Roche - workflow, data reductions/sparsification/SVD approximations for contractions
  - Andrei Alexandrou - overlap analysis campaigns for KNL-s, gauge operators/spin physics

- Thermodynamics and Workflow (co-PI: Swagato Mukherjee)
  - Sergey Panitkin: PanDA/ATLAS workflow for LCF systems, multi-site campaigns, scheduling, file transfers & data integrity
  - Subcontract: Alexei Bazavov - transport coefficients

Achievements so far

- Gauge generation (clover sea quarks): now 73 times faster for Summit;
  - seeking ways to bring this achievement to HEP codes (CPS, MILC).
- First demonstrations of multi-site job submissions under new “Harverster” job manager within PanDA
- Code generation technology from QUARC picked up by OpenMP as demonstration for next gen layout/zone execution
- Tensor contractions implemented in TACO/Tiramisu framework

Recent Highlight - sent to ASCR/NP, soon OLCF/JLab highlight

Accelerating QCD Gauge Generation on GPUs

Objectives

- Discovery of the properties of hadronic and nuclear matter through world leading Lattice Quantum Chromodynamics (LQCD) calculations
- Extension of the state of the art in LQCD computational capability by the development and integration of advanced algorithms
- Maximally exploiting advanced leadership hardware capabilities such as GPUs in OLCF Titan and OLCF Summit

Impact

- Nearly 2 orders of magnitude efficiency increase for gauge generation using OLCF Summit
- Titan improvements bring nearly an order of magnitude increase in value from existing INCITE allocation for USQCD gauge generation program
- Improvements fundamentally shift the balance of costs between gauge generation and gauge field analysis, allowing previously unaffordable calculations

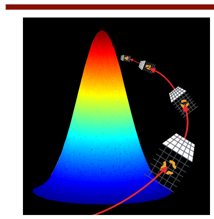

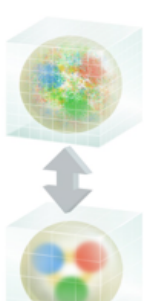

Image Credit: Joanna Griffin, Jefferson Lab Public Affairs

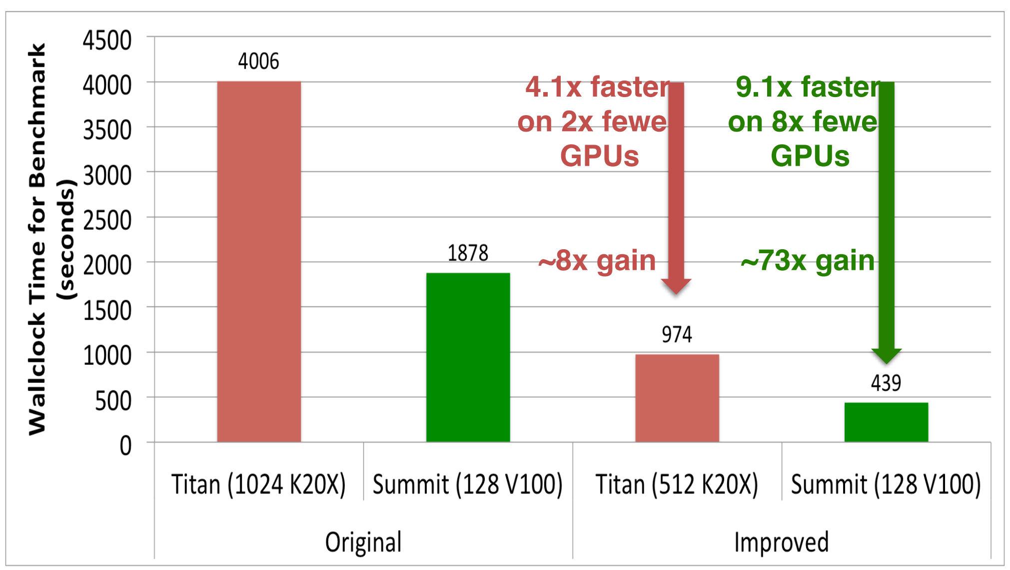

Accomplishments

- **~9x** wallclock speed-up on Summit using **8x** fewer GPUs than Titan: **~73x improvement in computational efficiency**
- Moved Multi-grid set-up phase in QUDA library entirely to GPUs allowing its use in gauge generation, and added extra optimizations (K. Clark, NVIDIA)
- Multi-grid solver integrated into Gauge Generation code (B. Joo, JLab)
- Developed Force-Gradient Time-stepper for Chroma (B. Yoon, LANL)
- Re-tuned Hamiltonian splitting and multi-level integration scheme enabled by these advances (B. Joo, JLab)

Computing Properties of Matter with Leadership Computing Resources

U.S. DEPARTMENT OF ENERGY | Office of Science | Jefferson Lab Exploring the Nature of Matter | Los Alamos NATIONAL LABORATORY EST. 1943 | NVIDIA

Tensors are everywhere

Science and Engineering

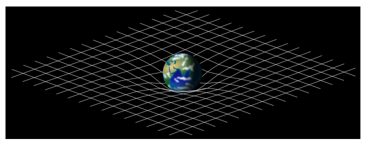

General Relativity

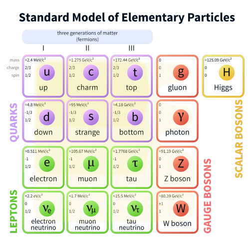

QCD

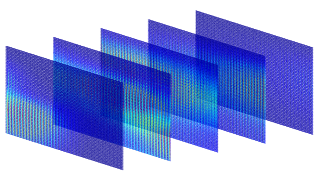

Fluorescence spectroscopy

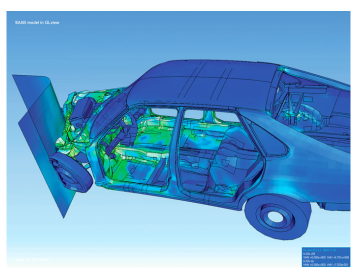

Finite Element Method

Data Analytics

Movie ratings

Product Reviews

Social interactions

Machine Learning

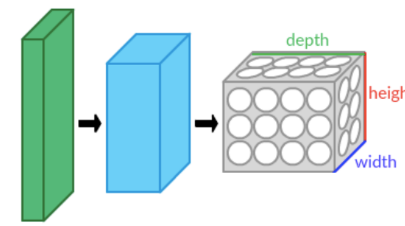

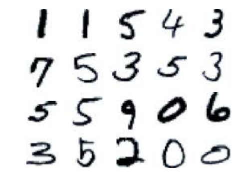

Images

Convolutional Layer

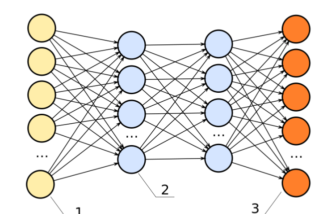

Connected Layer

The Tensor Algebra Compiler (taco)

$$A_{ij} = \sum_k B_{ijk} c_k$$

$A : (\text{dense}, \text{dense})$
$B : (\text{dense}, \text{sparse}, \text{sparse})$
$c : (\text{sparse})$

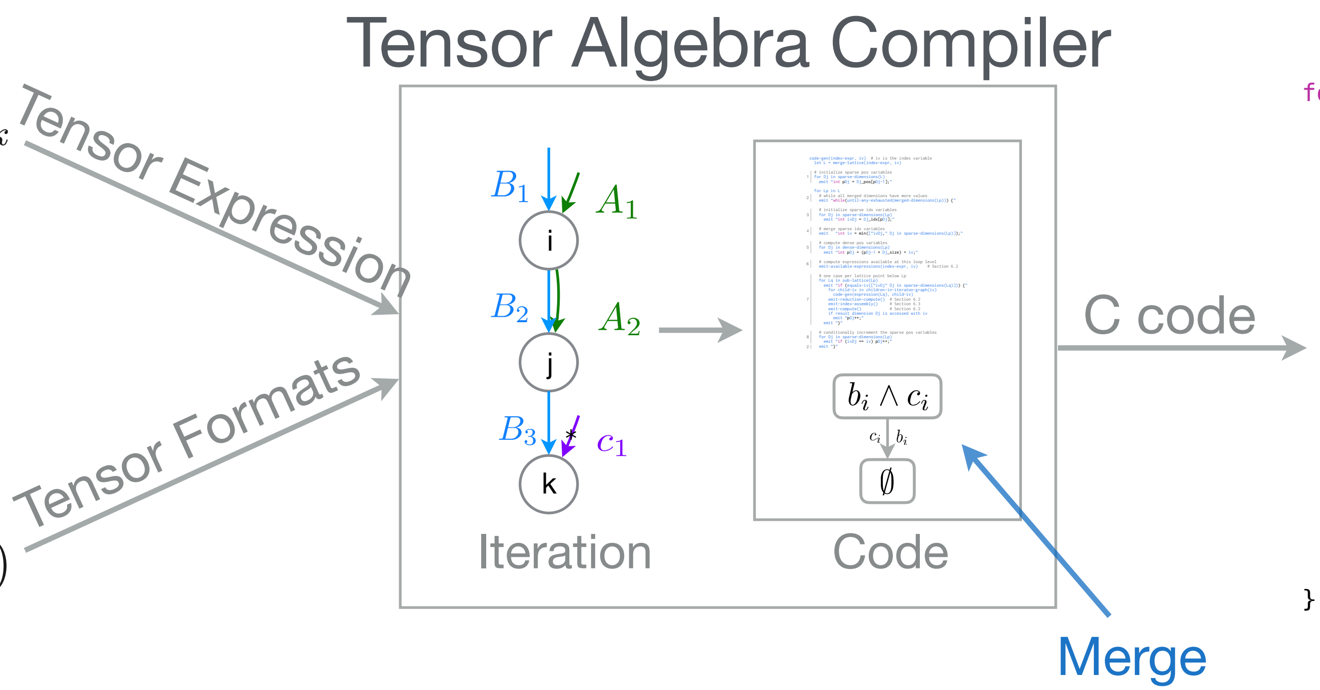


```
for (int i = 0; i < m; i++) {

  for (int pB2 = B2_pos[pB1]; pB2 < B2_pos[pB1 + 1]; pB2
    int j = B2_idx[pB2];
    int pA2 = (i * n) + j;

    int pB3 = B3_pos[pB2];
    int pc1 = c1_pos[0];
    while (pB3 < B3_pos[pB2 + 1] && pc1 < c1_pos[1]) {
      int kB = B3_idx[pB3];
      int kc = c1_idx[pc1];
      int k = min(kB, kc);
      if (kB == k && kc == k) {
        a[pA2] += b[pB3] * c[pc1];
      }
      if (kB == k) pB3++;
      if (kc == k) pc1++;
    }
  }
}
```

Fredrik Kjolstad, Shoaib Kamil, Stephen Chou, David Lugato, and Saman Amarasinghe

Heterogeneous Computing for Halide Using Tiramisu

MIT: **R. Baghdadi**, S. Amarasinghe, M. Ben Romdhan, J. Ray, E. Del Sozzo, Google: P. Suriana, Adobe: S. Kamil

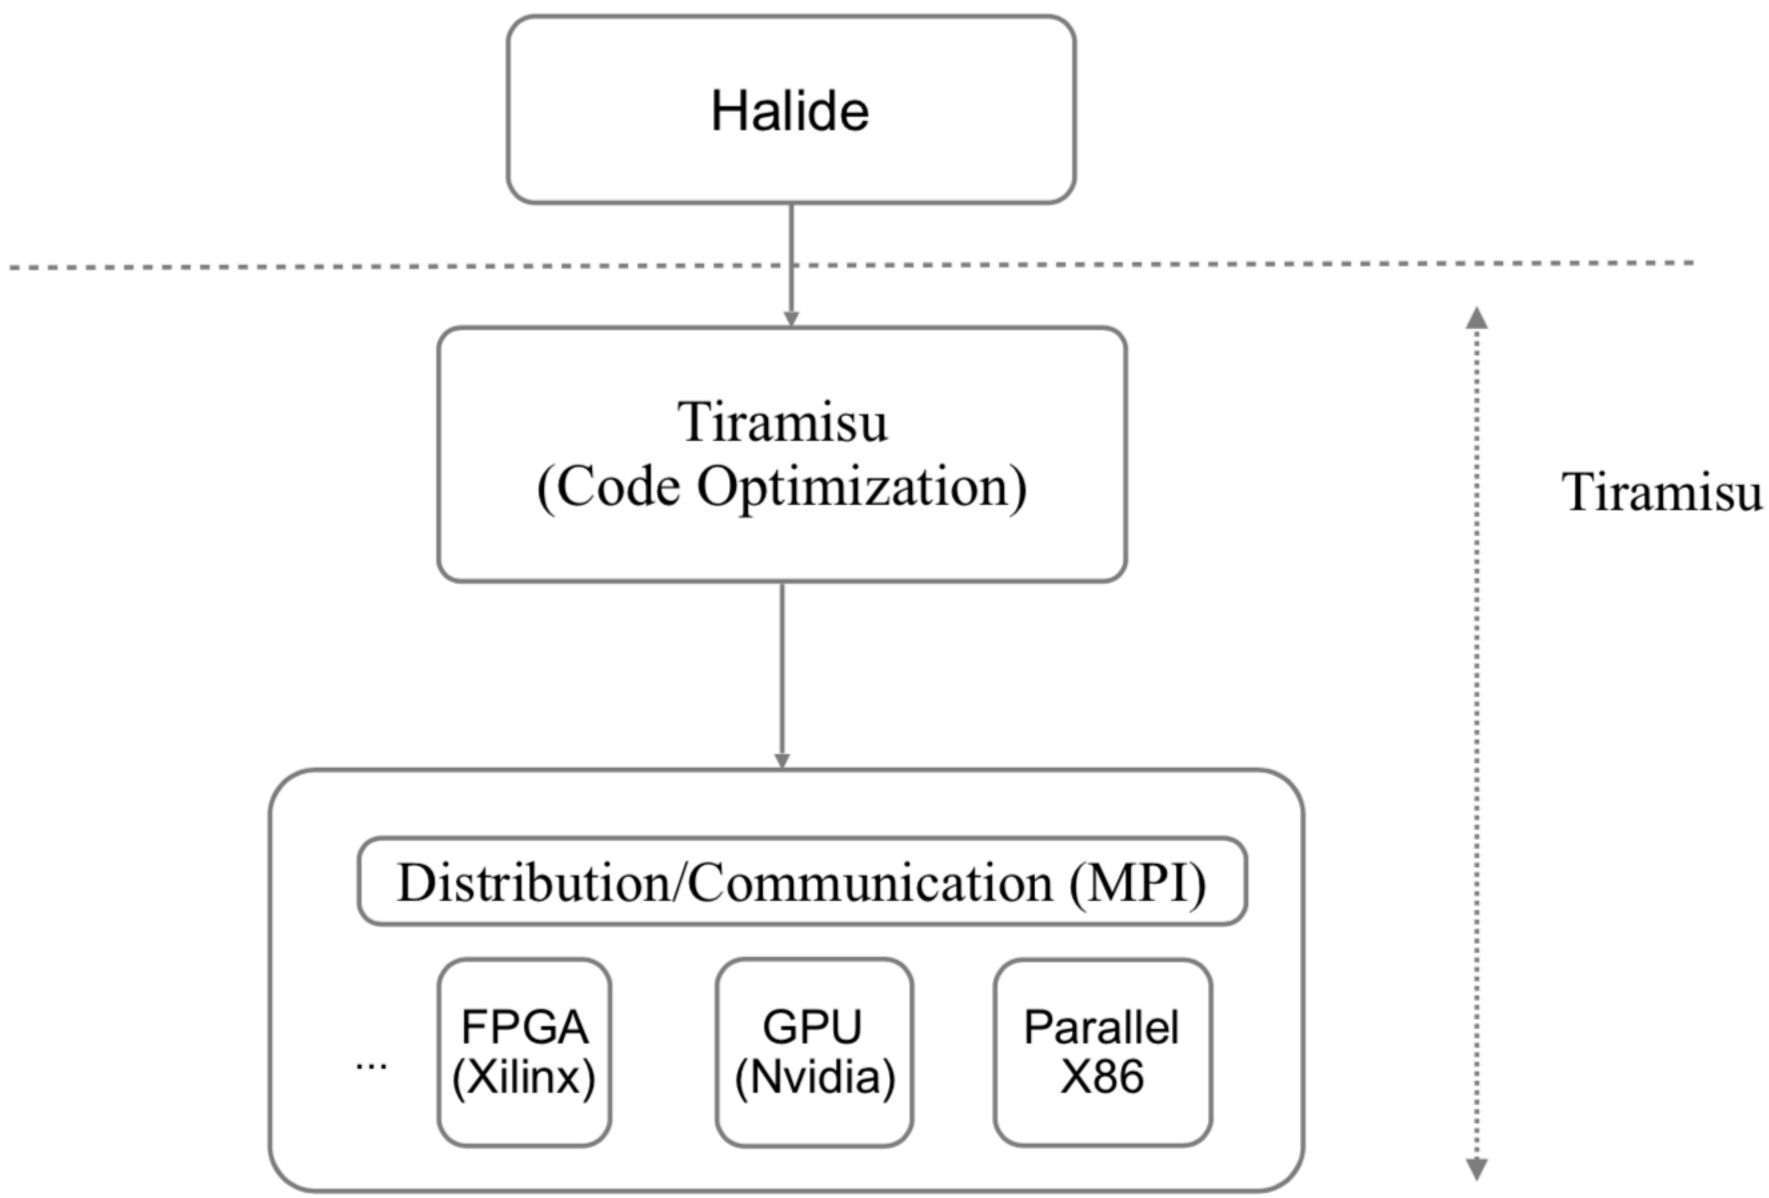

PanDA Workload Management System

Sergey Panitkin and Pavlo Svirin
for
BigPanDA project

PanDA in ATLAS

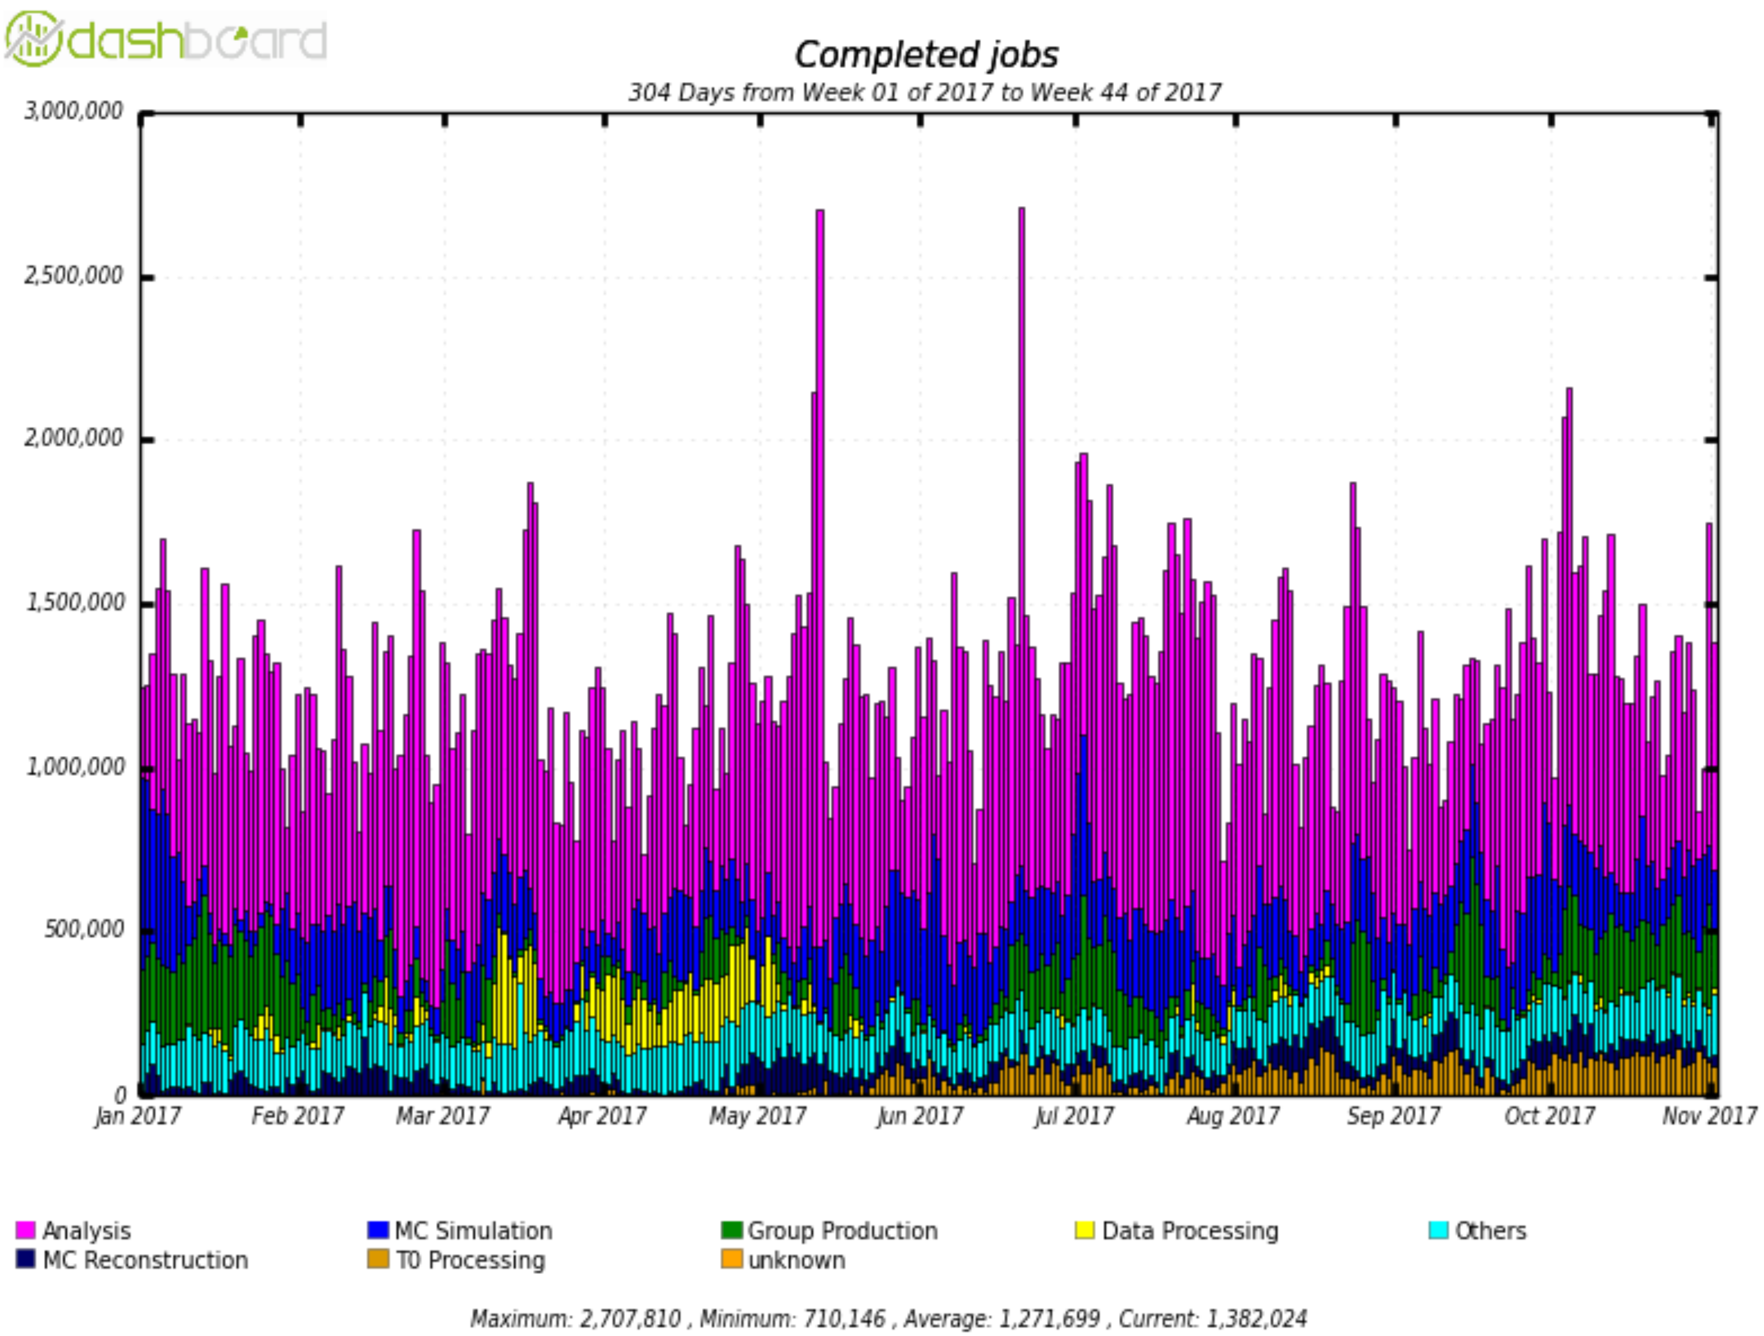

All ATLAS workloads are managed by PanDA
Up to 2.7M completed jobs per day, ~1.3M jobs per day on average

WLCG MoU Signatures

2017:
- 63 MoU's
- 167 sites; 42 countries

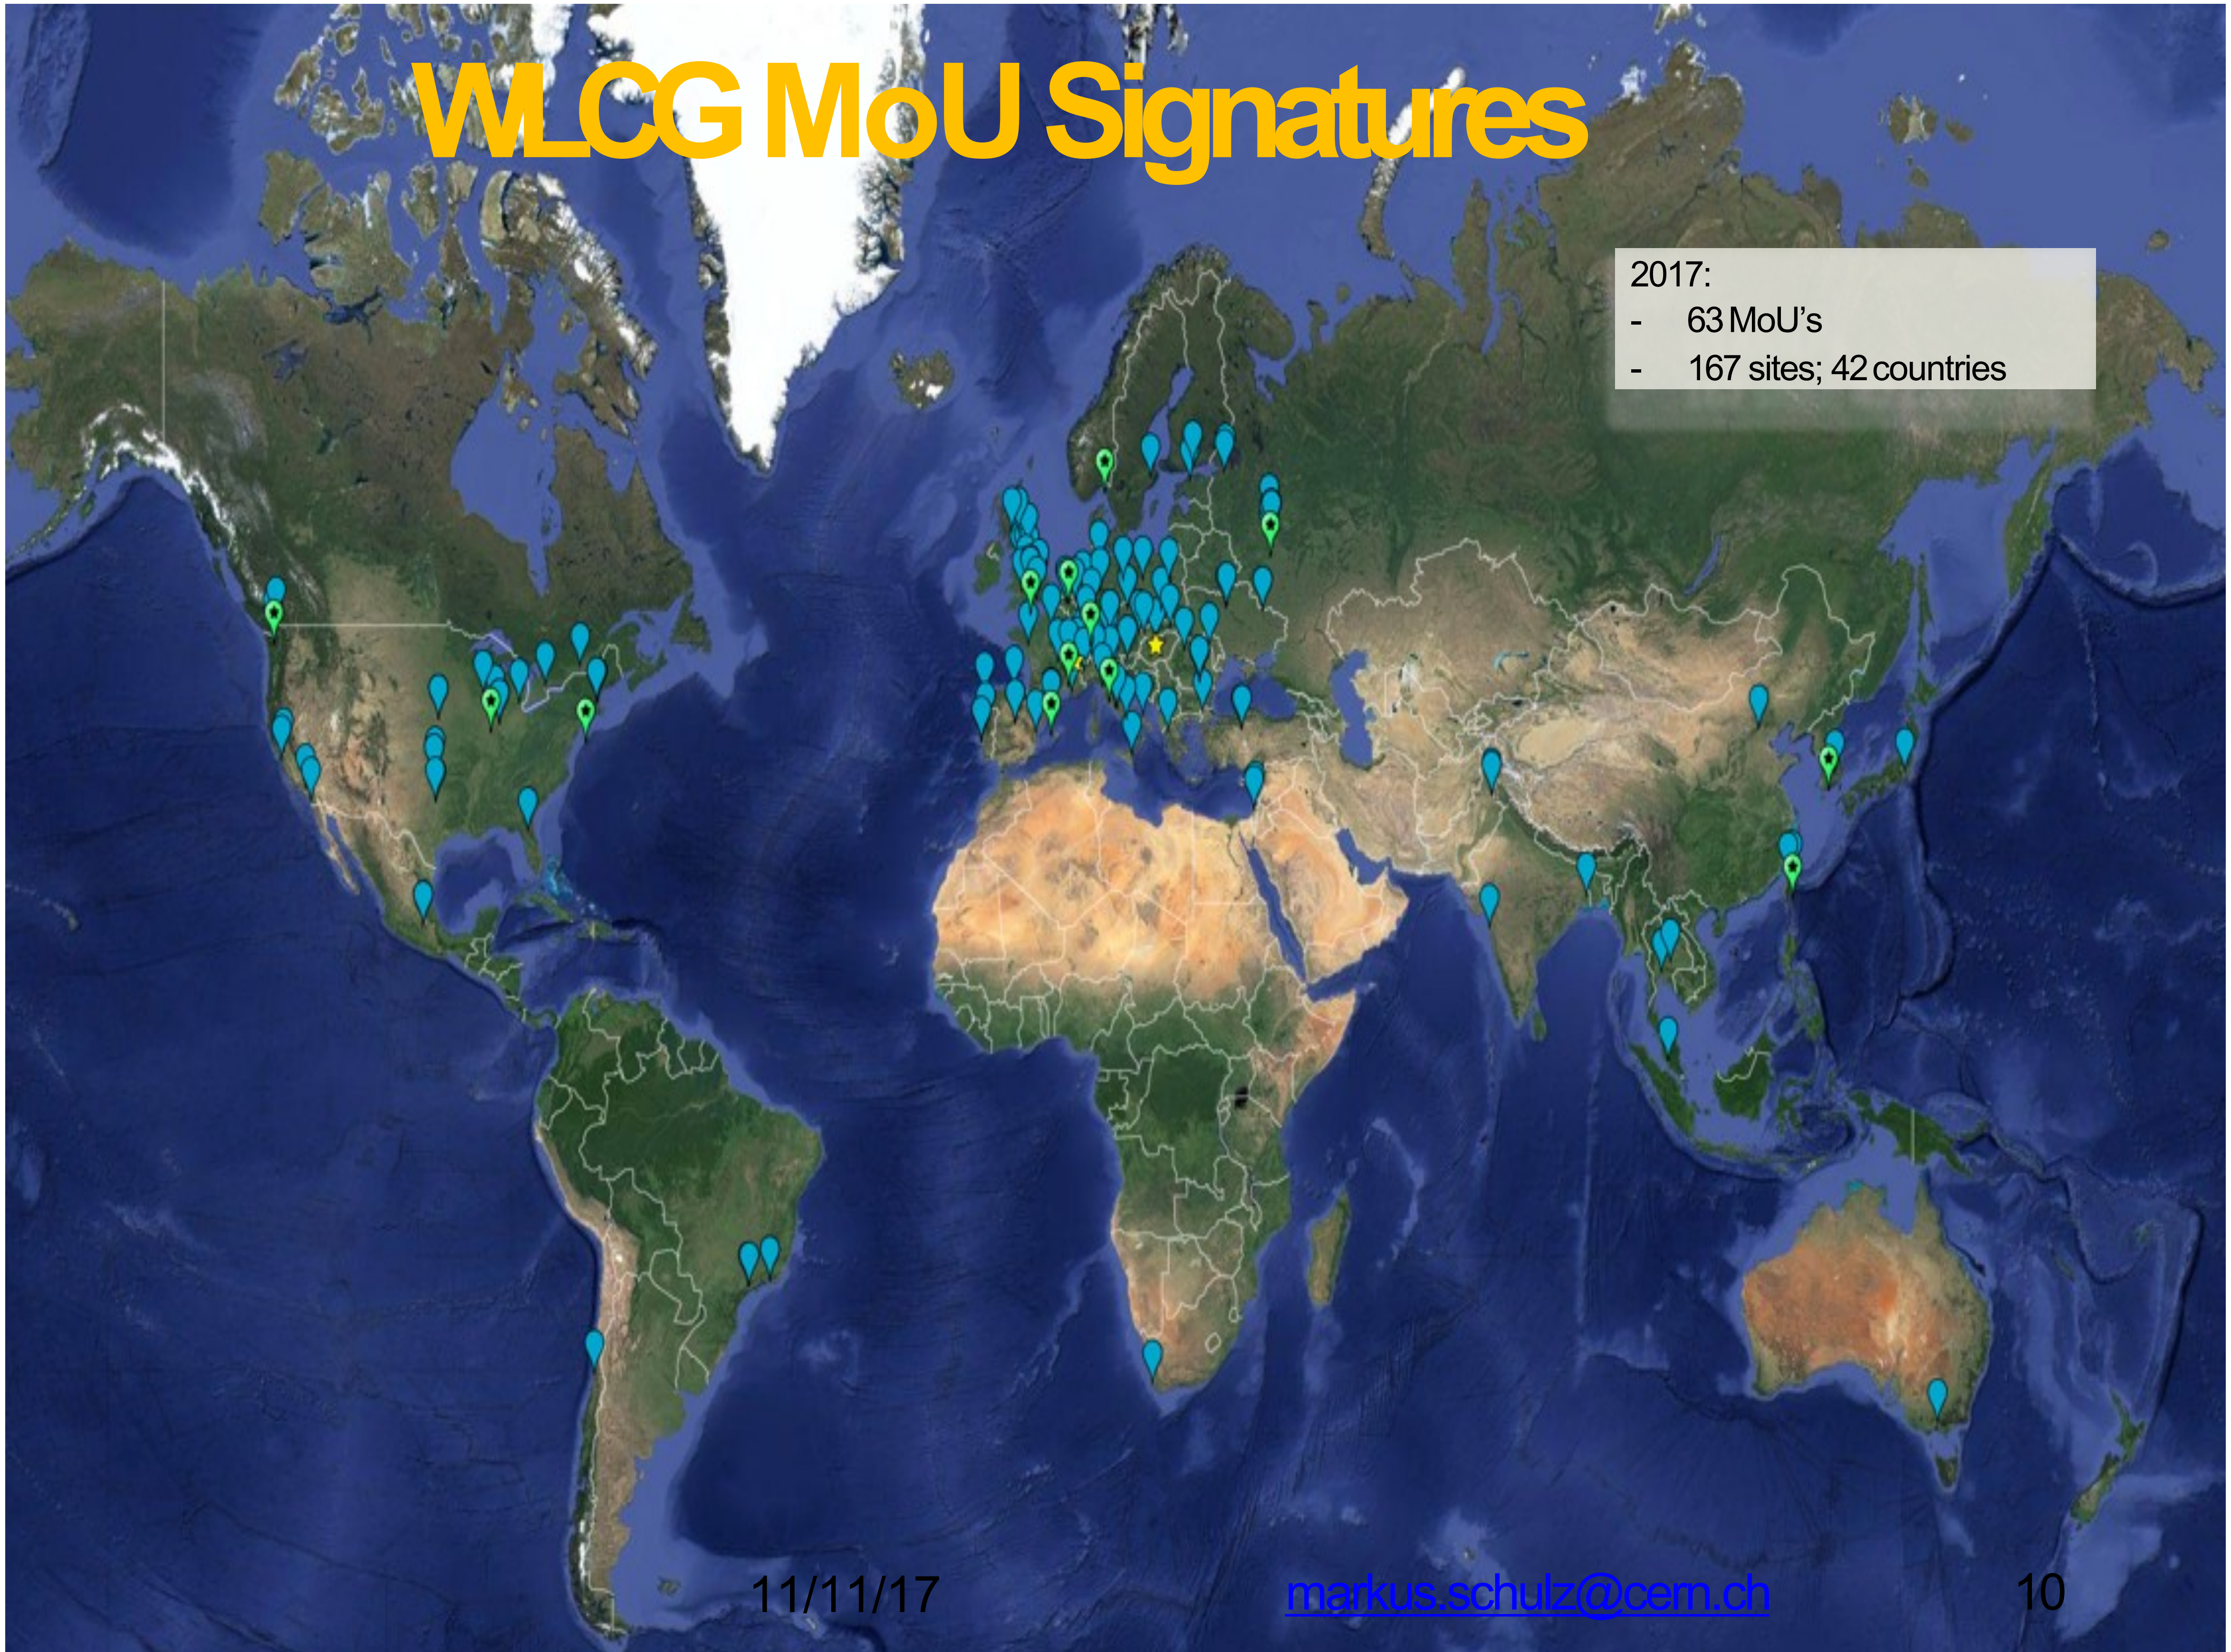

PanDA Workload Management System

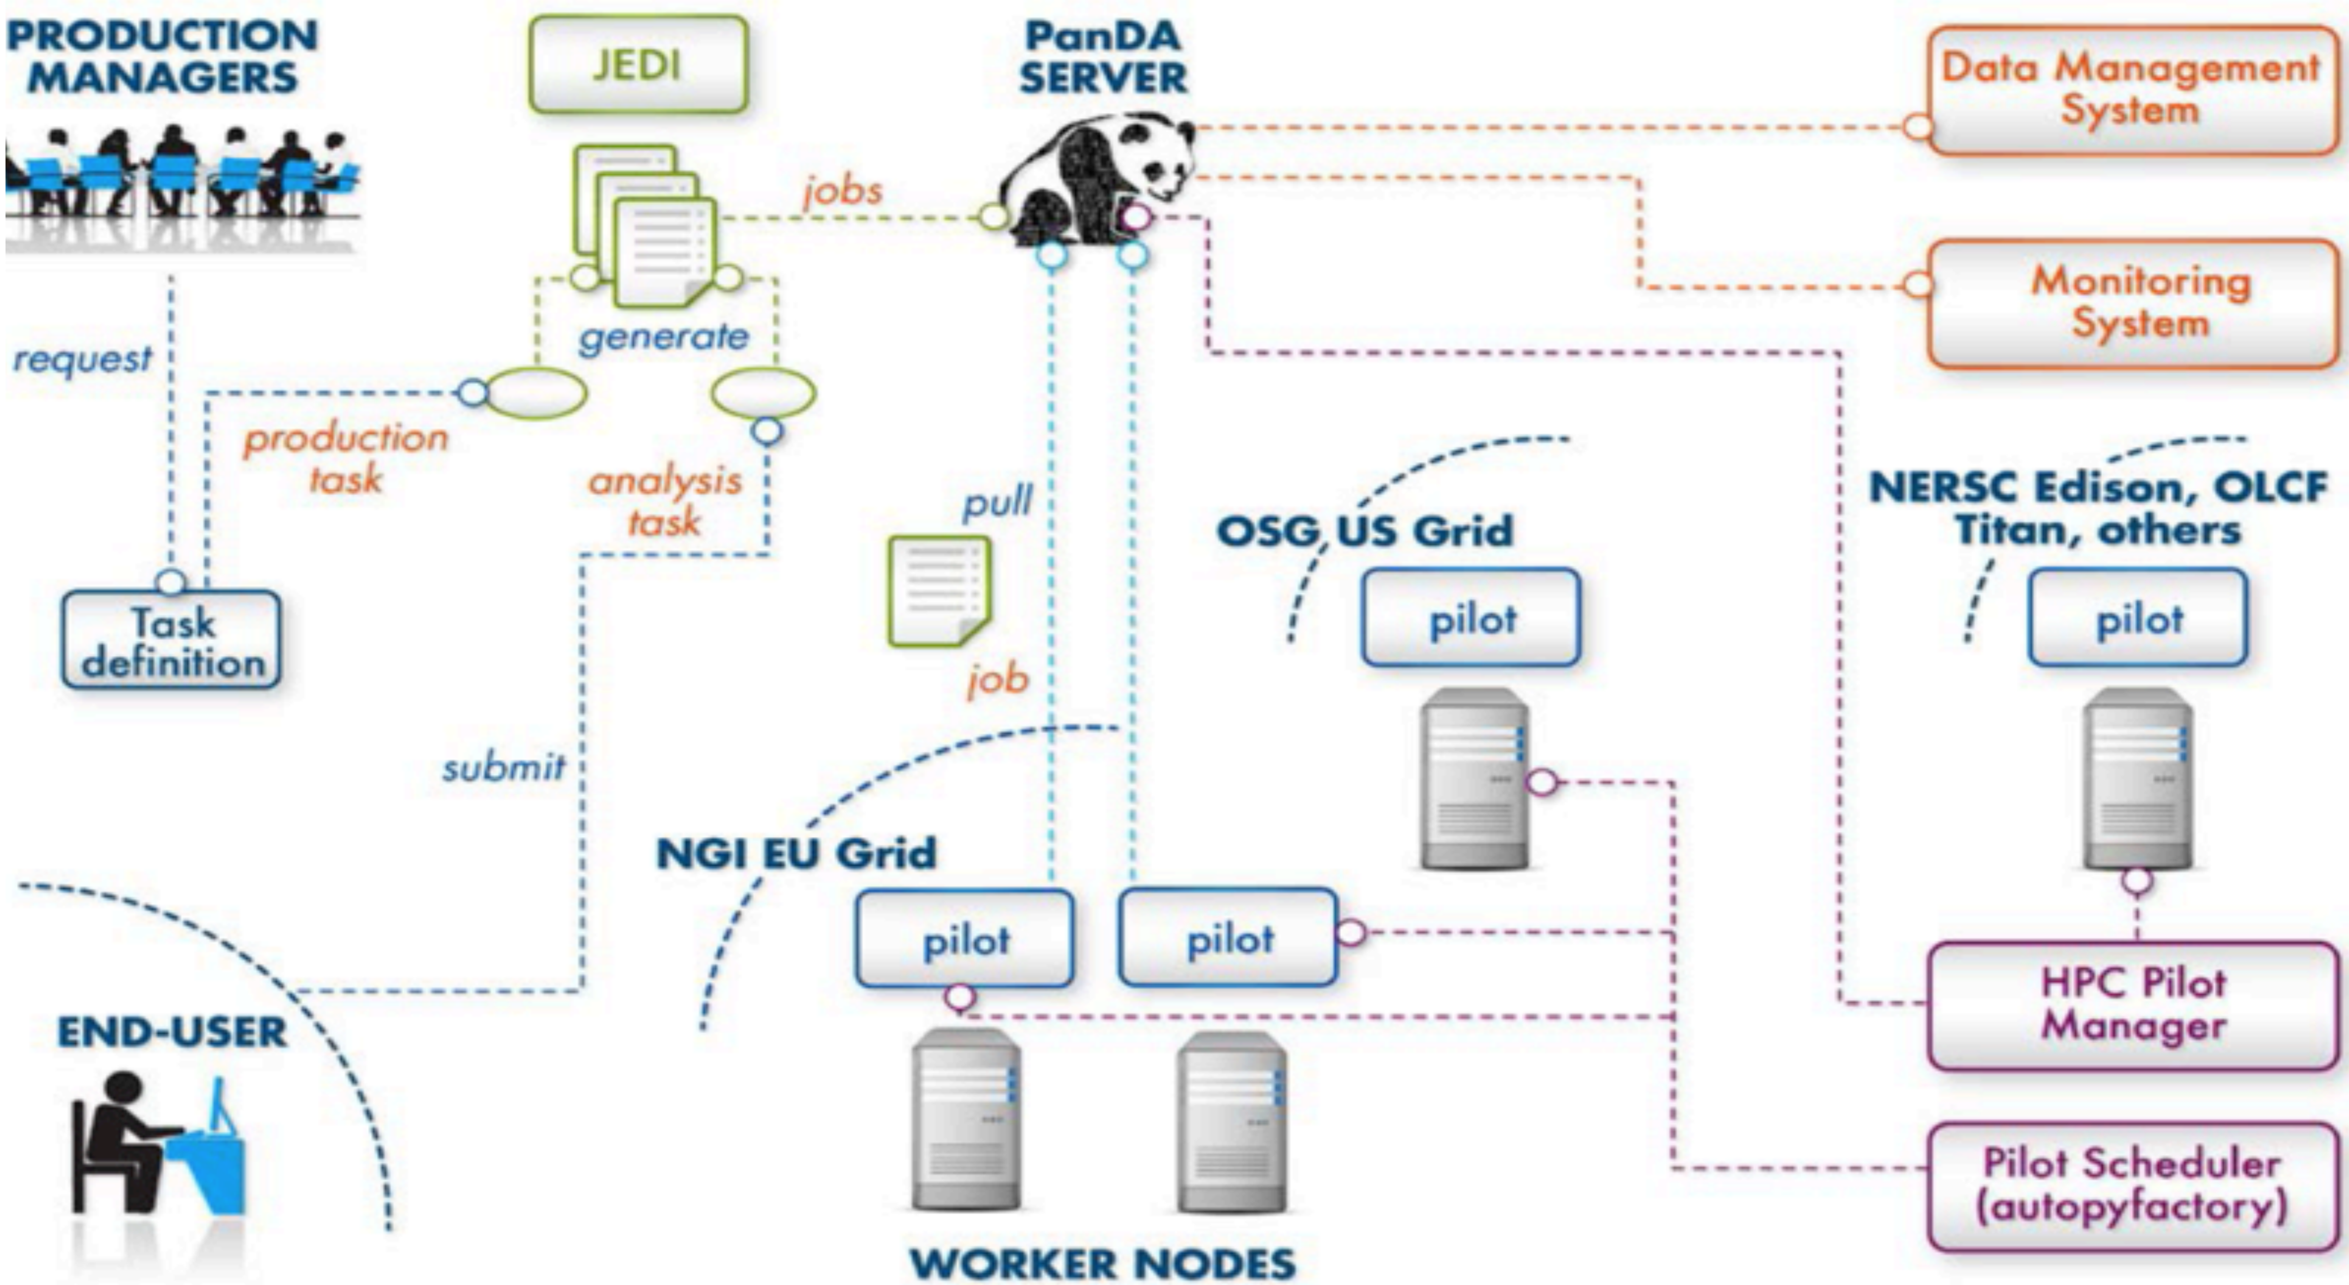

ATLAS operations on Titan

- Consumed ~170M Titan core hours from 1/17 to 1/18

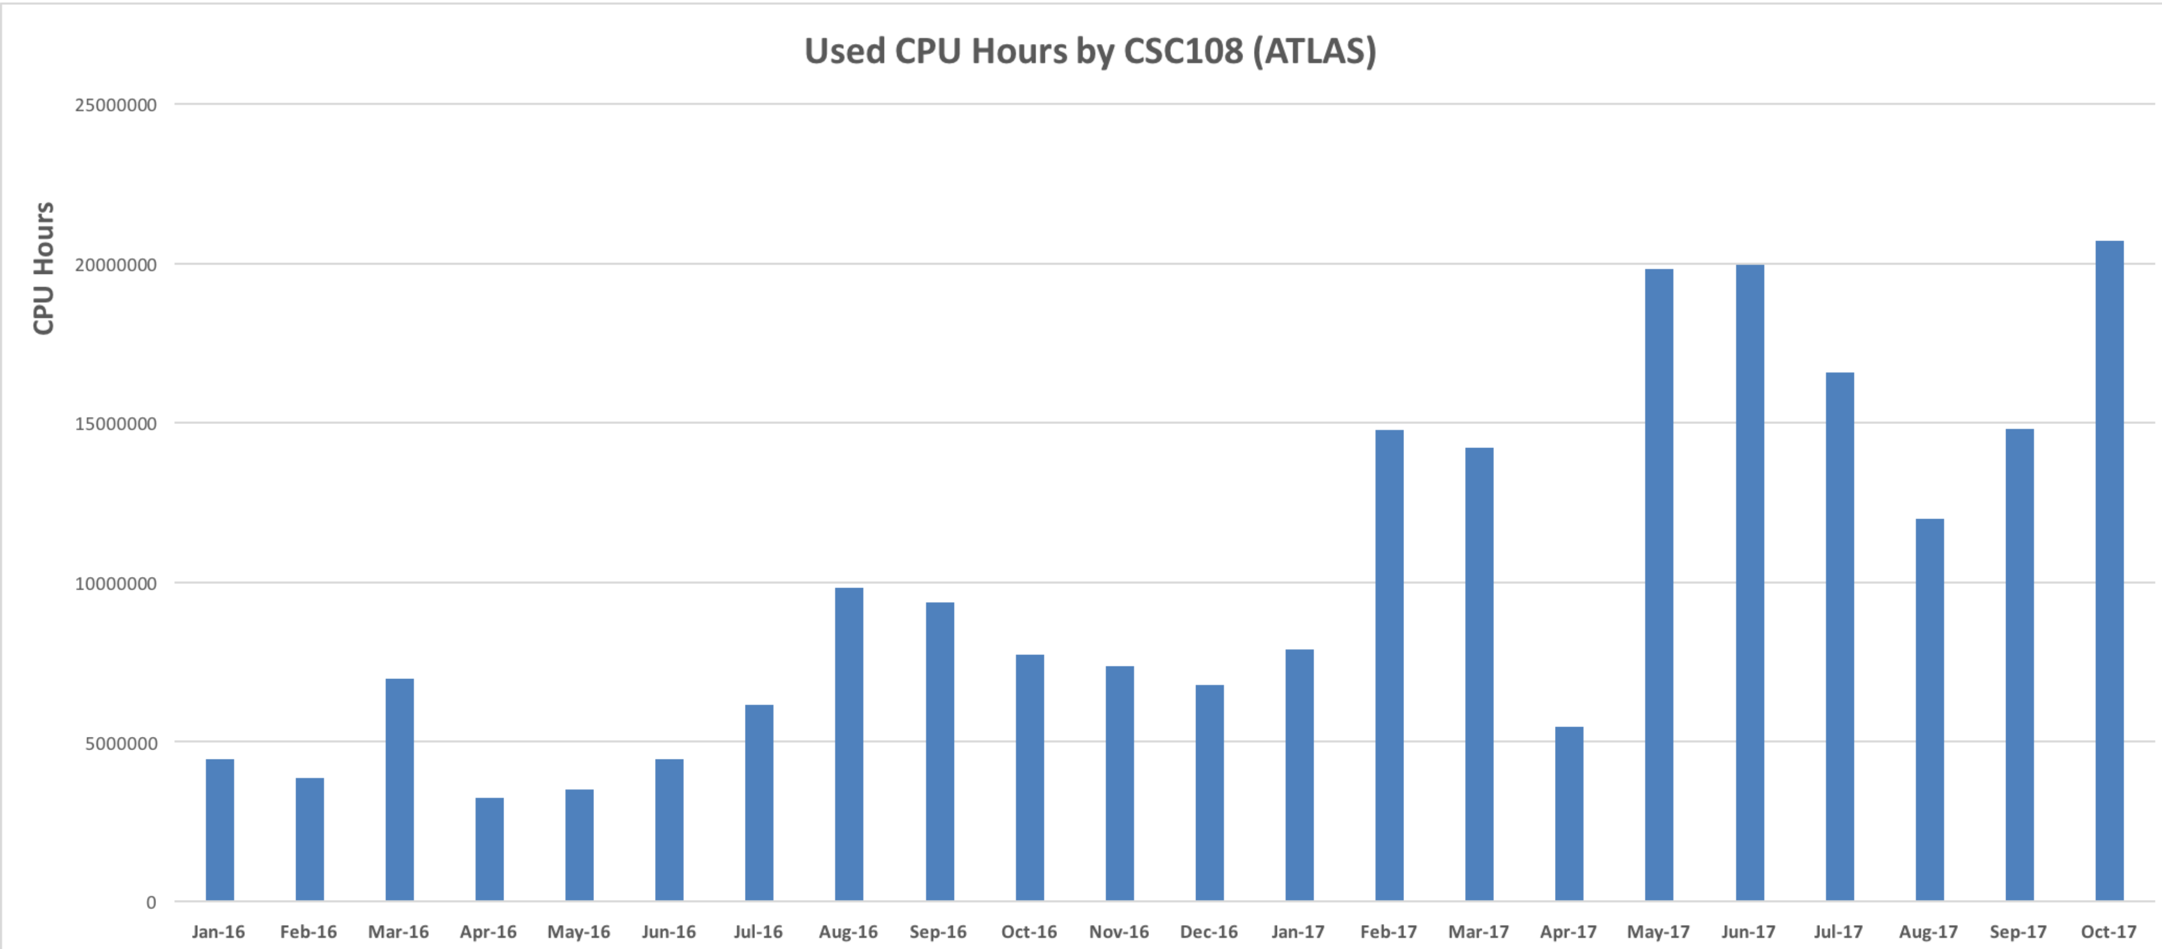

Job size shaping helps to achieve short wait times on Titan ~ 70 seconds

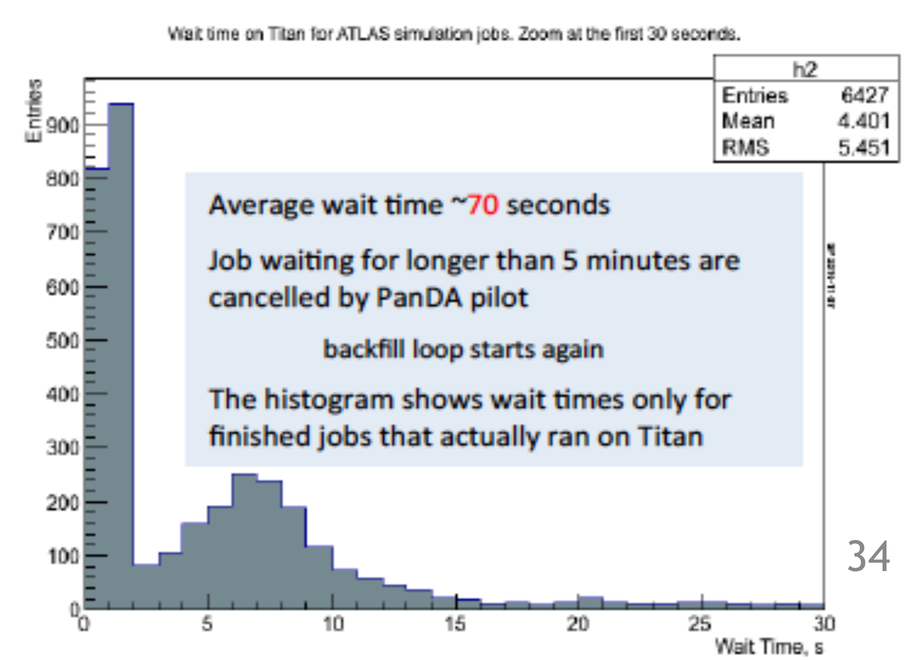

Exascale Computing Project

USQCD/ECP Overview

- FY17 budget \$2.5 M per year. FY18 likely to be about the same.
- Develop/prepare algorithms and software for lattice QCD calculations on the planned Exascale systems – first one in 2021.
- 50X improvement in scientific output – quantified with a set of benchmark problems.

Personnel

Robert Edwards	JLab	Barbara Chapman	SUNY Stony Brook
* Bálint Joó	JLab	Lingda Li, PhD student	SUNY Stony Brook
* Frank Winter	JLab	Andreas Kronfeld	Fermilab
* Jie Chen	JLab	Paul Mackenzie	Fermilab
* Daniel Trewartha, postdoc	JLab	* A. Strelchenko	Fermilab
Andreas Stathopoulos	William and Mary	Carleton DeTar	U. Utah
* Eloy Romero, PhD student	William and Mary	* A. Vaquero, Postdoc	U. Utah
* Chulwoo Jung	BNL	Steve Gottlieb	Indiana U
* Meifeng Lin	BNL	* Yuzhi Liu, postdoc	Indiana U
* Martin Kong	BNL	Bill Gropp	U. Illinois
* Yong-Chull Jang	BNL	* Paul Eller, PhD student	U. Illinois
Bob Mawhinney	Columbia U	James Osborn	Argonne NL
* Jiqun Tu PhD student	Columbia U	* Xiao-Yong Jin, postdoc	Argonne NL
Norman Christ	Columbia U	Richard Brower	Boston U
* Yidi Zhao PhD student	Columbia U	* E. Weinberg, postdoc	Boston U
Peter Boyle	Edinburgh U		
Guido Cossu	Edinburgh U		

* At least some salary support from ECP

Organization

- Solvers: Rich Brower (BU)
  - Boston U, Illinois, William and Mary, BNL
  - Develop communications avoiding algorithms. Multigrid. Pipeline CG, etc.
- Critical slowing down mitigation: Norman Christ (Columbia)
  - Columbia U, ANL, Edinburgh, BNL
  - Identify new molecular dynamics algorithms that accelerate diffusion of topological charge
- Contractions and matrix elements: Robert Edwards (JLab)
  - JLab
  - Optimize the evaluation of many thousands of quark-line graphs needed for nuclear matrix elements.
- Software: C.DeTar (U Utah)
  - ANL, Columbia, BNL, JLab, Edinburgh, Indiana U, BU, FNAL, U Utah
  - Develop software for supporting the above activities. Optimize codes.
- OpenMP development: B. Chapman
  - SUNY SB, BNL

Responses to 2017 Report

Suggestion #2

- Given the growth of young researchers in the field, the collaboration should consider adding additional junior members to its executive and scientific program committees. The new directions for the project proposed in item 1 above suggest USQCD should consider the election of a new spokesperson and new personnel in its executive and science policy committees.

- *Response: The Scientific Program Committee has had a larger fraction of younger members for several years. For example, one of the new members added this year is junior faculty. In response to this suggestion, this year's changes to the Executive Committee brought in members significantly younger than those they replaced. We are also staging the second election for a junior member. As a federation of science collaborations, the USQCD charter calls for the new spokesperson to be selected by the Executive Committee, and this process was followed this year.*

Suggestion #3

- Data sharing (configurations) is part of the collaboration's charter. However, a data management plan was not presented at the review. The USQCD collaboration should develop such a plan and disseminate it at its All Hands meeting.
- <u>Response</u>: *This is a good suggestion. We have appointed a committee, headed by Deputy Spokesperson Robert Edwards, to develop a data-management plan. It will be posted on the USQCD website, so that members can use it as a foundation for their own data-management plans.*
- (Mentioned in Robert's talk.)

Suggestion #4

- Since physics deliverables are the ultimate objective of the project, the definition and documentation of science milestones should be paramount. The project should develop procedures to document scientific milestones uniformly over all the LQCD areas so that the project can track their annual progress quantitatively.

- *Response: We agree with this suggestion. We have been trying for several years to make the presentation of the goals in our various sub-fields more concrete and more uniform, and believe that we have made progress. We plan to make milestones clear in the 2018 whitepapers.*

- (Note also that too many rigid milestones would prevent researchers from reacting to events.)

Suggestion #5

- Given the direct relevance of lattice gauge calculations to the experimental community, it would be valuable to enlist experimental physicists to advocate for the project during future reviews and/or the next multi-year extension proposal past 2019.

- Response: *This is an interesting suggestion. We are well aware that the support of experimenters is crucial to our success. How to include them in a review is tricky because we have several diverse subject areas that would each require different speakers, so a balanced presentation from experimenters at a review would consume a lot of the review. How to marshal the support of experimenters in proposing the extension of our funding is very important. We will investigate how best to obtain the advice and support of experimental physicists as we move forward.*

- (Letters of support of a FY20-24 Project?)

Suggestion #6

- The feedback from the User Survey indicates a high user satisfaction with the project and its allocation process. The project is encouraged to continue taking such surveys. One suggestion to improve feedback to the project is to hold a user-organized session during the annual All Hands Meeting to discuss user perspectives of the allocation process and the facility operations. Such a session, if actually user motivated and well attended, may be a good way to more clearly capture any common user pain points for using the facilities.

- *Response: Every All Hands' Meeting has reports from all three sites managers with time for questions, and the Project Manager's report always presents the User Survey. It is common for questions, suggestions, and complaints to be aired at this time. For example, in the 2018 meeting, a discussion of BNL operations (which are now different with the IC model) came up. It was very constructive, giving the site manager lots of useful information to share with colleagues. Other features, both good and bad, at the other sites were also discussed. Thus, while we share the sentiment behind this suggestion, we believe the usual agenda accomplishes these aims, and has for some time.*

Suggestion #7

- If the project moves to Institutional Clusters as the main provider of cycles in its capacity computing model, then USQCD should consider the election of a new spokesperson and new personnel in its executive and science policy committees to reflect this new approach.

- <u>Response</u>: *As discussed in the response to Suggestion #2, USQCD now has a new spokesperson, the new role of deputy spokesperson, and an overall younger Executive Committee. It may be worth noting that BNL worked closely with LQCD ext. II and with USQCD to design and procure hardware in a way similar to our previous designs of dedicated hardware. It should be emphasized that the program funds which the DOE invests in an institutional cluster must be spent wisely on hardware that will be highly cost effective for and address the computational needs of USQCD. Insuring this outcome remains one of the important functions of USQCD leadership.*

Suggestion #8

- USQCD will prepare a proposal for hardware purchases beyond FY19. USQCD should seriously consider institutional clusters. The project's rationale for purchasing their own hardware made more sense when they were first adopters of new architectures. This position is no longer true. An option for an FY19+ proposal could be to request funds to equip institutional clusters with features that may not be purchased otherwise, such as fast highly interconnected network systems. Such features would likely not harm non-LQCD users, but, as stated in the Future Plans presentation by P. Mackenzie, may be crucial for LQCD codes.

- Response: We agree with [most of] this suggestion. Through an existing allocation on the BNL Institutional Cluster, we began incorporating institutional clusters into our hardware portfolio in 2017. We are expanding our experience with institutional clusters in FY18 by collaborating with BNL on the acquisition of a new cluster configured to meet LQCD computing needs. In FY18, USQCD will run on new and existing IC hardware at BNL. New LQCD hardware at Fermilab, if procured in FY18, will follow the IC model, paying close attention to experience gained through the BNL acquisition process l. The Office of Nuclear Physics, on the other hand, prefers the dedicated-hardware model at JLab, where some advantages of the IC model have been part of operations for a few years. It may be useful to point out that in both the institution-cluster and dedicated-hardware models USQCD continues to be an aggressive first adopter of new hardware with important benefits to the larger HPC community. While our highly effective early adoption of GPUs is now history, the recent purchase of KNL machines at JLab and as part of the IC at BNL reduced costs by the early use of single- (JLab) and dual-rail (BNL) Omnipath networks. With collaborators in Edinburgh and Intel, USQCD solved highly technical difficulties that this network posed, ultimately determining the direction adopted by Intel to make this offering competitive [arXiv:1711.04883].

USQCD Leadership in Science:
Planned and Serendipitous

Planned Leadership

- Lattice Meets Experiment: quark flavor, BSM; neutrino-nucleus scattering, computing, ….
- Muon g–2 Theory Initiative: spearheaded by El-Khadra* and Lehner, brings everyone together. *as Fermilab Distinguished Scholar
- Beam Energy Scan I and II: planned with leadership from BNL lattice-QCD theorists
- GlueX experimental program planned with JLab lattice-QCD theorists.
- $|V_{ub}|$ and $|V_{cb}|$ from Λ_b decays: lattice-QCD [arXiv:1503.0142] worked closely with LHCb leading to 2015 Nature Physics paper.
- Neutrino Theory Network (ASK on Steering Board).

Serendipitous Leadership

- Discovery of light 0^{++} in near-conformal theories spurred interest from experimenters (e.g., Arce) and phenomenologists (e.g., Romanino, Bai).
- Discussions with Fermilab neutrino experimenters led to USQCD involvement in NuSTEC whitepaper.
- Quark-mass project led to better theoretical understanding of quark-mass definitions in QFT [arXiv:1701.00347, arXiv:1712.04983] ⇒ top quark.
- Discovery of smooth crossover changed view of early universe.
- Nuclear theory now considers lattice QCD as central to the field.
- Main simulation algorithm (pre-LQCD), hybrid Monte Carlo (HMC), now used as MCMC in many fields, known as Hamiltonian Monte Carlo (HMC).