U.S. DEPARTMENT OF ENERGY | Office of Science

Future HEP Computing Needs
(DOE Perspective)

May 31, 2019

Eric R. Colby

Senior Technical Advisor, Office of High Energy Physics

Office of Science, U.S. Department of Energy

Outline

- DOE-HEP’s efforts to understand the challenge
  - HCI-WG
  - Community Roundtable Meetings
  - Code Analysis Program
  - ECP efforts for HEP
- Next Steps

Motivation

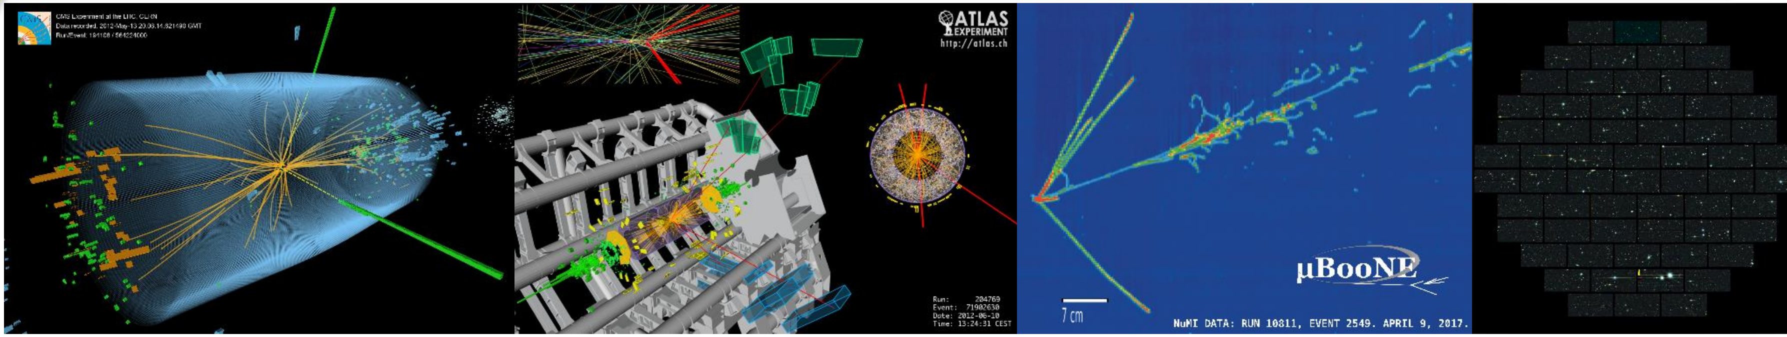

- HL-LHC, DUNE, and LSST will produce complex data at rates an order of magnitude or more beyond current experiments
- Successful implementation of the broad science program envisioned by P5 will require an equally broad and foresighted approach to the computing challenges
  - P5 assumed that data analysis could occur from anywhere on the planet. HEP computing resource pressures threaten to break this assumption.
  - Meeting these challenges will require us to work together to more effectively share resources (hardware and software) and appropriately integrate commercial computing and HPC advances.
- Last year's data survey of the three HEP frontiers showed that "business as usual" is likely to cost ~7 times the current computing budget
- Earlier discussions with the community revealed (2017) a pervasive belief that "business as usual" would be adequate. **This has been changing.**

HEP Computing Infrastructure Working Group

Stood up in 2017, the HCI-WG is a working group of OHEP personnel charged with:

- Developing and maintaining an HEP Computing Resource Management Strategy, and
- Recommending actions to implement the strategy.

The HCI-WG has:

- Conducted an initial survey of the computing needs from each of the three physics Frontiers, and assembled a preliminary model.
  - The preliminary model showed that computing costs will exceed expected funding by a **large factor** within the next decade
- Assisted with allocations of NERSC resources and provided OHEP's recommendations concerning ALCC proposals
- Pursued a consultative process with the HEP community to (1) more accurately capture the largest expected computing needs and (2) look for opportunities where economies of scale and optimal use of resources can close the gap

HEP Computing Model (May 2018)

- In preparation for the Inventory Roundtable, the largest HEP experiments from all three frontiers were asked to provide a **more detailed estimate** of their expected computing needs
  - CPU, storage, network, personnel, and HPC portability
- Cost estimates for all experimental frontiers:
  - "Business as usual" (minimal additional HPC use): **$600M ± 150M**
  - With effective use of HPC resources this reduces to: **$275M ± 70M**
- By 2030 cost share by frontier is estimated to be:
  - ½ Energy Frontier
  - ¼ Intensity Frontier
  - ¼ Cosmic Frontier
- **A strategy encompassing all HEP computing needs is required!**
  - **Heterogeneous need (not just LHC), and cost drivers (not just CPU)**

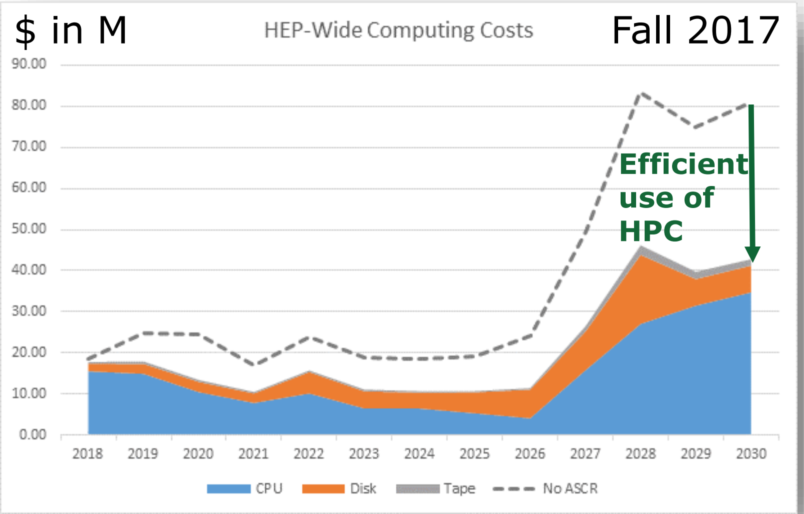

Courtesy: Tom LeCompte

HEP Computing Strategy Development

- **We must develop a better strategy to address the computing challenges ahead, or miss out on science**
- **Recognizing this, a consultative process was begun in 2018:**
  - **Inventory of HEP Computing Needs Roundtable Meeting (May 2018)**
    - Focused on hardware, software, and personnel needs for the next decade
    - Made clear the need to dive into a spanning set of codes to better estimate the cost of using GPU-accelerated architectures, and the estimate the potential benefit
  - **Code analysis projects**
    - CCE* charged August 2018 to assist with connecting code experts with LCF software experts to do deep-dive analysis. Supplementary funding was offered to support these activities.
    - FWPs are just now coming in.
  - **Commonalities Roundtable Meeting** – Fall/Winter of FY 2020
    - Will focus on identifying common elements in software and workflows, HPC applicability, and integration with Exascale, HSF, S2I2, and other computing initiatives
    - Identification of recurrent kernels and themes in HEP software
  - **Future community roadmapping exercise?**

*CCE = HEP Center for Computational Excellence, see https://press3.mcs.anl.gov/hepfce/

HEP Inventory of Computing Needs Roundtable
(May 7-8, 2018)

Department of Energy
Office of Science
Washington, DC 20585

Dr. Ben Brown, SC-ASCR
Via e-mail: Benjamin.Brown@science.doe.gov

March 26, 2018

RE: HEP Inventory of Computing Needs Roundtable Meeting

Dear Ben,

The Office of High Energy Physics (HEP) of the U.S. Department of Energy (DOE) is developing a strategic plan for the computing hardware, software, and workforce needed to address the significant computing challenges inherent in the 2014 P5 strategy for particle physics.

Analysis of the input gathered in initial surveys conducted by HEP in 2017 indicated that computing costs will exceed available funding by a large factor within the next decade. Developing an alternative plan, or possibly an evolution to the current plan, is critical, and requires a more detailed look at the anticipated computing needs over the next decade, focusing on areas where economies of scale and optimal use of resources can close the gaps.

We invite you to participate in the development of a strategy to address this gap. We are charging a small focus group composed of representatives of the largest national laboratory computing directorates, major HEP experiments, and major computational efforts to provide HEP with summaries of anticipated computing hardware, software and tools, and personnel needs over the next decade, and share these summaries at the HEP Inventory of Computing Needs Roundtable Meeting, a 1 ½ day meeting to be convened in the Washington D.C. area in the first two weeks of May.

Successful implementation of the broad science program envisioned by P5 will require an equally broad and foresighted approach to the computing challenges inherent in conducting a program that delivers on the science for discovery. We hope you will accept the invitation and help DOE formulate a strategy to address these challenges.

If you will accept this invitation, please email your response and calendar availability for the date range May 1st through May 11th to Thomas.LeCompte@science.doe.gov .

Sincerely,

James Siegrist
Associate Director
for High Energy Physics

HEP Inventory of Computing Needs Roundtable Meeting
May 7-8, 2018
Cambria Hotel & Suites, Rockville, MD

Room: Sinequa A/B

	Day 1		
8:30 AM	Welcome, Motivation, Objectives for the Meeting	Jim Siegrist	0:20
8:50 AM	The Challenges Ahead	Tom LeCompte	0:40
9:30 AM	HPC Examples: HACC	Salman Habib	0:15
9:45 AM	HPC Examples: WarpX	Jean-Luc Vey	0:15
10:00 AM	HPC Examples: LQCD	Ruth Van de Water	0:15
10:15 AM	*Break*		0:15
Lab-by-Lab Overviews of Computing Resources and Strategy			
10:30 AM	ANL & ALCF+OLCF	Salman Habib	0:55
11:25 AM	Brookhaven	Kirsten Kleese van Dam	0:35
12:00 PM	*LUNCH*		1:50
1:50 PM	Fermilab	Liz Sexton-Kennedy	0:45
2:35 PM	LBNL & NERSC+ESNet	Peter Nugent	0:55
3:30 PM	SLAC	Richard Dubois	0:25
3:55 PM	*Break*		0:15
4:10 PM	Summary of Needs from Data Call	Tom LeCompte	0:35
4:45 PM	Discussion	ALL	0:45
5:30 PM	Adjourn		

Room: Sinequa A/B

	Day 2		
Experiment-by-Experiment Overviews of Computing Needs			
8:30 AM	Objectives for the Second Day	Tom LeCompte	0:15
8:45 AM	ATLAS	Paolo Calafiura	0:30
9:15 AM	CMS	Oliver Gutsche	0:30
9:45 AM	IF Experiments at Fermilab	Panagiotis Spentzouris	0:30
10:15 AM	*Break*		0:15
10:30 AM	LSST/DESC	Katrin Heitmann	0:20
10:50 AM	DESI	Peter Nugent	0:15
11:05 AM	Cosmological Simulations and Archiving (incl. CMB-S4)	Salman Habib	0:25
11:30 AM	Closing, Next steps	Jim Siegrist	0:15
11:45 AM	Discussion	ALL	0:45
12:30 PM	Adjourn		

What We Learned from the First Roundtable

- While HPC architectures will continue to evolve, moving to vectorized, multithreaded codes tailored to I/O-bound systems will result in higher efficiency codes
  - Engaging HPC experts to analyze code has helped identify algorithm alternatives and data flow bottlenecks, in some cases resulting in spectacular speedups (e.g. 600x). Continued engagement is therefore essential!
  - Need to identify which codes could benefit the most
- Using Exascale machines badly (e.g. by ignoring the GPU/accelerator) will result in a factor-of-40 penalty in performance that will not be tolerated. HEP will lose its allocations if it does this.
  - Engaging Exascale Computing Project (ECP) experts early and often will result in faster adoption of best practices for exascale machines, and influence ECP design choices to HEP's benefit. HEP needs a coordinated interface to both ECP & the Leadership Computing Facilities.
  - Need to identify which codes could benefit the most
- LQCD regularly rewrites its code, has ***reaped significant speedup benefits*** every time
- Reinforced that multiyear NERSC allocations & better metrics for pledges are needed
- End-to-end network data flow models are needed to support tradeoff analysis of storage vs. CPU vs. network bandwidth on a system-wide and program-wide basis
  - Greater sharing of the underlying data management software layer may also be beneficial

CCE Charge: Code Analysis

August 15, 2018

Charge

HEP-CCE is charged with (1) finalizing with the Office of HEP the list of selected software packages for this initial round of analysis [see below], (2) contacting the Points of Contact (POCs) for each software package and facilitating interaction with appropriate HPC experts at NERSC and the LCFs, (3) tailoring, then negotiating the scope of work for each case.

List of selected software packages

Criteria for selection:

- Widespread use by the HEP community now and in the future
- A diversity of types, providing a broadly representative sample of the implementation challenges to be expected of HEP codes in general

The Office of HEP proposes the following codes and POCs:

- ***GEANT4***
- ***Track following and vertexing*** codes for the ATLAS and CMS Collaborations, analyzed **together**
- ***Sherpa***
- ***GalSim***
- ***LArSoft***
- The Office of HEP welcomes discussion with HEP-CCE to refine this list

Facilitating interaction with appropriate HPC experts

- HEP-CCE is asked to consider the specific analysis needs of each code and pair the code POC with appropriate LCF experts, facilitating the contact as needed.
- Funds will be provided to support this activity. The HEP-CCE is asked to balance and coordinate FWP requests to HEP once code POC/LCF expert assignments are known.
- Report to the Office of HEP an estimate as to when the analysis of all examples in the list will be sufficiently advanced to permit an informed discussion at the Commonalities Meeting.

Tailoring and negotiating the scope of work

This exercise is focused on three objectives

- Assess the current state of the code
  - Analysis of the structure of the code
  - Identify prioritized list of resource-intensive sections of the code
- Identify high-impact HPC-compatible upgrades
  - Identify opportunities for speedup at the local level (including I/O issues)
  - Identify opportunities to better exploit HPC libraries and "motifs" being developed by the Exascale Computing Project
  - Identify the basic features needed in a proxy code
- Estimate the time and effort required to implement such upgrades

Plans for the “Commonalities” Roundtable

Fall/Winter 2019

Second HEP Computing Roundtable

Identification of Opportunities for Synergy

~~July~~ August 2018

v. 3~~1~~ 2/28~~14~~/2018

Desired Outcomes – Commonalities

- An appraisal of the overlaps ~~and~~ across, relevance of, and gaps between the specific recommendations in the HSF CWP and S2I2 Strategy to the US HEP computing strategy
- Identification of examples of portions of major codes that could be reworked to boost performance to test both HPC speedup and a ‘skunkworks’ approach to software modernization.
  - Identification of key bottleneck locations in major data analysis workflows
  - Identify curriculum for the Code writer’s school
- Ideas on how to foster closer community collaborations; shared sense of overall problems in HPC/HEP computing
- Strategy and map for working together on reconstruction software for HL. How will I/O and storage problems in general be addressed?
- Discuss workflow models.
- Reinforce education of thought leaders in the computing community about the challenges ahead, and the process and technical approach that is being developed to address the challenges

Framing Questions

- [Already have LeCompte’s preliminary computing resource needs model]
- What are the opportunities for
  - efficient shared hardware usage,
  - stronger collaboration on software development and maintenance,
  - 12-month duration HPC code rework examples to boost performance,
  - more efficient adoption of industry products and practices?
- What national framework could effectively steer computing resources?
  - What advisory mechanism(s) could support a national computing resource management strategy?
  - What are the options for synchronizing and complementing the ECP effort?
  - What is/are the best mechanisms to (a) orchestrate the community and (b) communicate with the community?

Draft Agenda (8:00-3:00~~, working lunch, ~6 hrs usable~~)

- Context 90’
  - Second Warning Shot Across the Bow (in lieu of welcome) (Siegrist)

The timing of the Commonalities Roundtable will depend on the Code Analysis task, as this is a major input to the discussion

ECP* Efforts for HEP

"Projectized" Software R&D efforts to maximally leverage Exascale computing

- ~2.5M/year x 4 years for each effort

ExaSky – 100x Science Reach, 10x accuracy

- ***Goal:*** *100x increase in science reach, 10x increase in accuracy, contributing to a range of Cosmic Frontier Experiments*
- ***Physical Models****: Cosmological hydrodynamics coupled to gravity, (Collective) N-body particle transport, subgrid physics for multiple astrophysical mechanisms (heating, cooling, star formation, etc.)*
- ***Codes****: HACC (particle hydro+gravity), CosmoTools (analysis framework), Nyx (AMR hydro+gravity)*
- ***Motifs:*** *Monte Carlo, Particles, Sparse Linear Algebra, Spectral Methods, Unstructured Grid, Structured Grid, Graph Traversal*

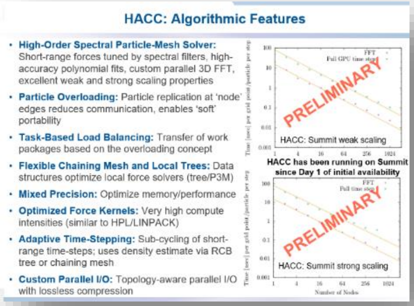

ExaSky PI: Salman Habib

LatticeQCD – 10x precision

- ***Goal:*** *10x increase in precision, uncover evidence of deeper particle structure*
- ***Physical Models:*** *Quantum Chromodynamics (the theory of quarks and gluons in the standard model), quantum electrodynamics.*
- ***Codes:*** *MILC, Chroma, CPS.*
- ***Motifs:*** *Sparse matrices, dense matrices, relaxation methods, Monte Carlo methods.*
- *MPI, OpenMP, CUDA, Kokkos, OpenACC, C++17, Thrust, SyCL, QUDA, QPhiX, LAPACK, ARPACK.*

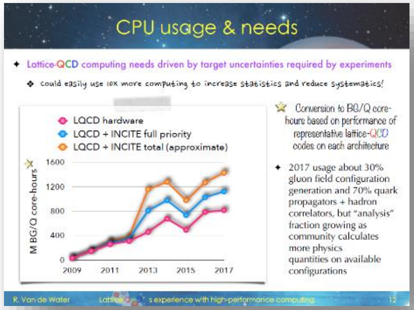

LatticeQCD PI: Andreas Kronfeld

WarpX – Full 1 TeV Plasma Wakefield Collider

- ***Goal:*** *Simulate 1 TeV plasma wakefield collider of 100 stages*
- ***Physical Models:*** *electromagnetic waves, charged particles dynamics, plasmas.*
- ***Codes:*** *Warp+PICSAR+Boxlib ➔ WarpX.*
- ***Motifs:*** *Particle-In-Cell, Maxwell solver (finite-difference and FFT-based), AMR on block structured grids.*

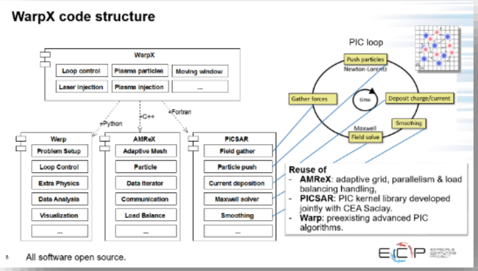

WarpX PI: Jean-Luc Vay

*ECP = Exascale Computing Project, see https://www.exascaleproject.org/

U.S. DEPARTMENT OF ENERGY | Office of Science

Code-Developer Schools and Traineeships

- The CMS experiment alone had ~100-150 programmers per month editing the code in the 2014-2016 time period
  - For all of HEP, the total number of active programmers is significantly larger
- Having a significant fraction of the active developer community conversant with best programming practices for next-generation hardware is critical
- Code-developer schools—training general HEP code developers in the best practices for next-generation architectures—will play a critical role
  - Preparing 50% of the CMS code developers by the end of Run 3 would mean ~70/year
- The sooner the schools start, the more science leverage HEP gains
- Preliminary discussion is underway about adding HEP-targeted HPC modules to (e.g.) the Argonne Training Program on Extreme-Scale Computing and other venues
- DOE Traineeships are another potential mechanism with a longer engagement time
  - Targeted at specific areas of need; require measured outcomes
  - Supports 2 years of tuition, stipend, and travel for trainees

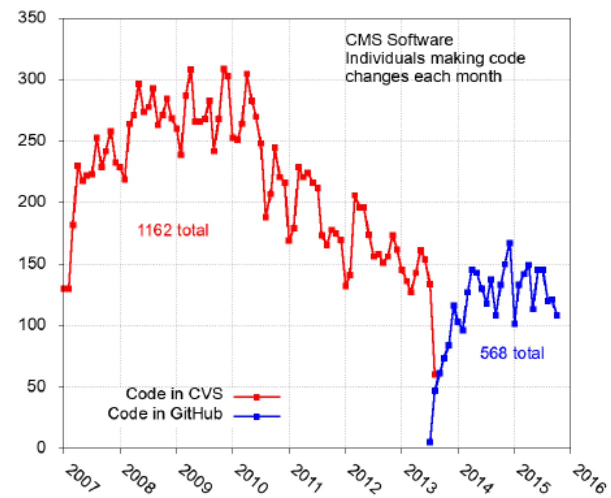

Figure 4 from *P. Elmer, M. Neubauer, M. Sokoloff, "Strategic Plan for a Scientific Software Innovation Institute (S2I2) for High Energy Physics", (2017).*

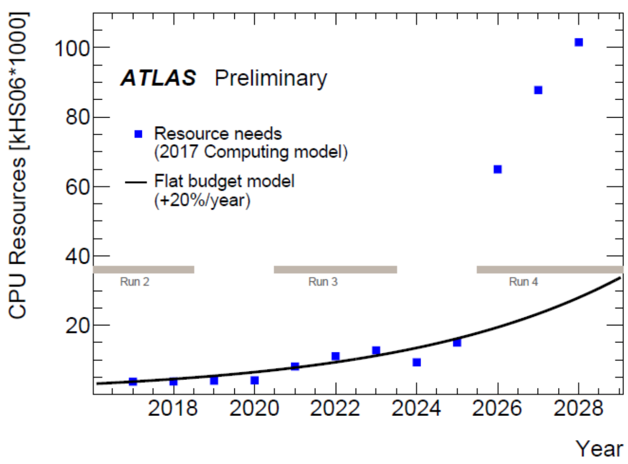

Figure 4a from *the HSF's "A Roadmap for HEP Software and Computing R&D for the 2020s", (2017).*

What’s Next

- DOE recognizes that Exascale computing will not solve all problems
  - DOE will focus on maximizing the scientific leverage that Exascale computing can deliver, working through the Labs
  - NSF’s IRIS-HEP is working on maximizing the scientific leverage that distributed high-throughput computing can deliver, working through the Universities*
  - DOE and NSF will continue to coordinate complementary approaches
  - Storage and network solutions are critical as data sets reach well beyond exabyte scale
- For DOE, CCE has played and will continue to play a key role facilitating use of HPC computing
- Conversion, verification, and validation of significant HEP codes is likely to require a significant, formalized effort
  - Perhaps similar to the Exascale Computing Project Co-Design Efforts
  - Drawing heavily on expertise in the ASCR community
- The first group of codes analyzed will guide DOE’s approach to this next step
- A “Roadmapping” exercise may be needed to identify goals, milestones, and develop a plan

* For a description of the IRIS-HEP program, see: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1836650

Additional Materials

One Paragraph Summary for the HEPAP Minutes

P5 recommended a program of challenging scientific experiments that have equally challenging computing needs. DOE-HEP stood up the HEP Computing Infrastructure Working Group (HCI-WG) in 2017 to develop a strategy for meeting the computing needs. The HCI-WG began a consultative process with the community, performing a data call in 2017, and the first of two roundtable meetings with HEP stakeholders. In addition to three ECP projects, a selection of six representative codes have been identified for further analysis for potential benefit (and cost) of making HPC implementations. A second roundtable will be held when the analysis is complete to assess the next steps. Developing and disseminating best programming practices for HPC, a community roadmapping exercise, and coordination with ASCR and NSF's IRIS-HEP will all be needed to successfully address the significant computing challenges ahead.

Why Won't Business as Usual Work?

- Moore's Law will not solve the problem

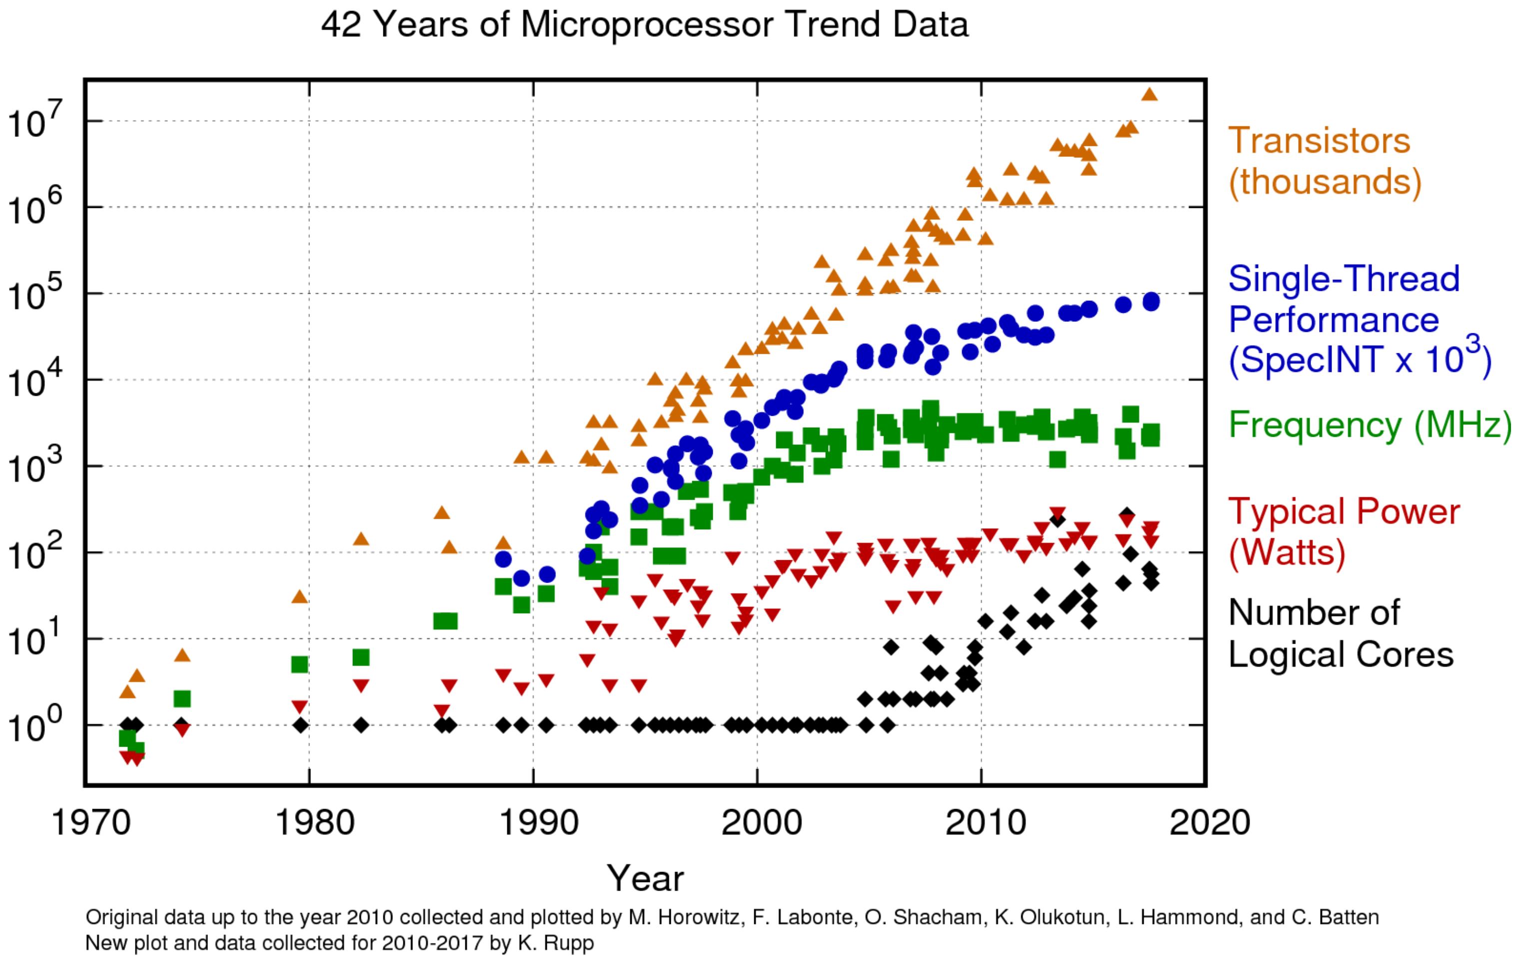

Why Won’t Business as Usual Work?

- Higher CPU node counts won’t necessarily solve the problem

Representative Plot of HEP Code Scaling
“Vanilla” Geant with idealized CMS geometry on Intel Phi

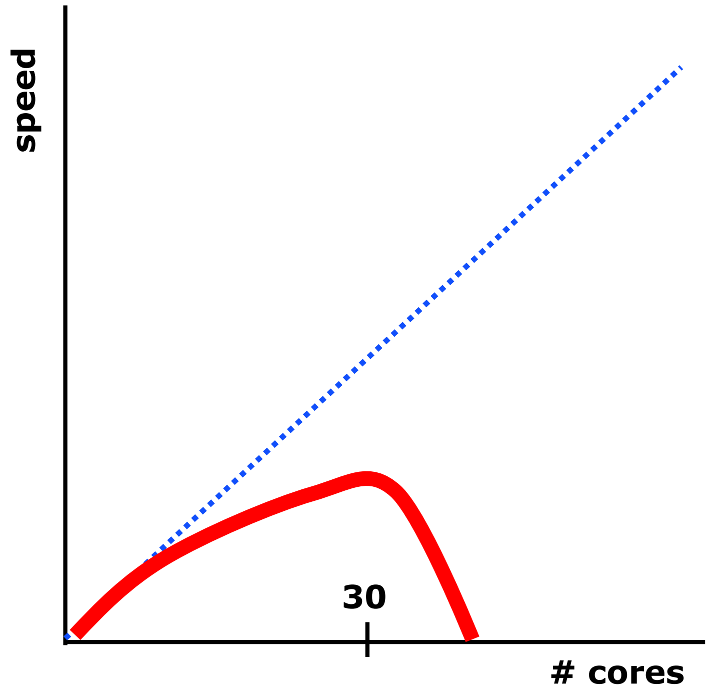

**** Thanks to Argonne’s Joint Laboratory for System Evaluation***

A Complicated Picture

Organizations Working on HEP Computing Issues

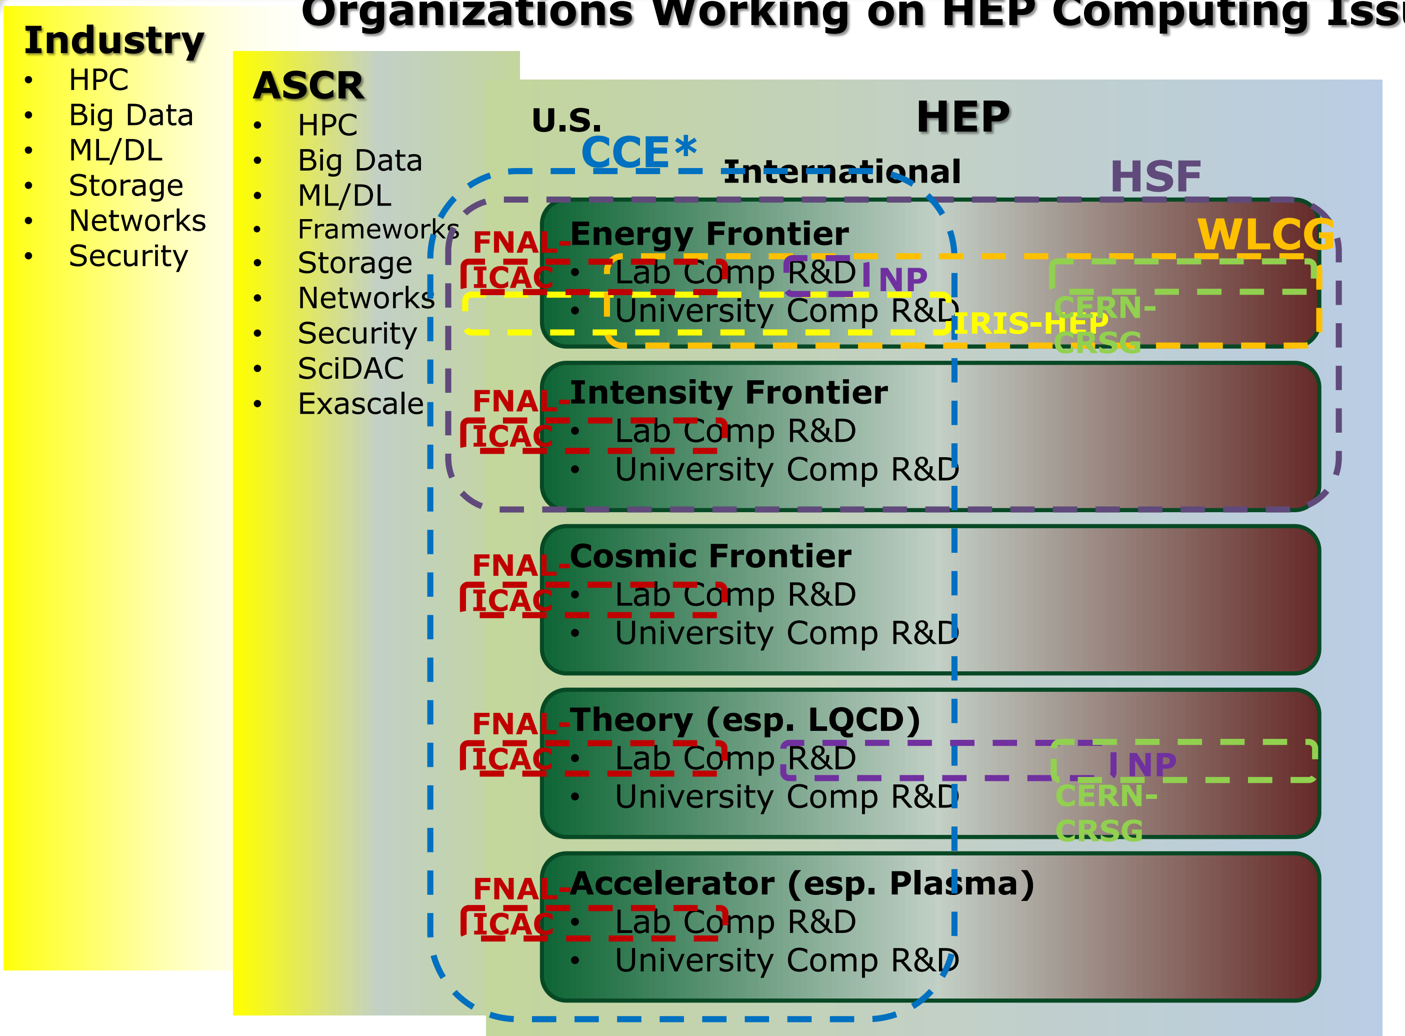

Other Active Orgs?
OSG, APS,
CERN Sci Forum Comp
EU-t0

*CCE focused on HPC issues, ML

U.S. DEPARTMENT OF ENERGY | Office of Science

Fermilab ICAC

Charter of the ICAC

- Reviews and Advises the laboratory on:
  - computing operations,
  - cyber security,
  - upgrade plans, and
  - software and computing R&D aimed towards
    - the development and exploitation of future facilities
    - as well as advancing scientific tools and methods in general
- Monitor progress with respect to the established laboratory objectives, currently encompassing:
  - Software and Computing for the Intensity Frontier Experiments;
  - Fermilab's involvement in the HL-LHC Software and Computing Upgrades;
  - Progress toward common solutions for the above domains;
  - National and International cooperation and collaboration with partner institutions;
- The ICAC is expected to address high-level strategic, programmatic, and planning issues, rather than specific implementation details.

Fermilab

2 14/03/2019 Liz Sexton-Kennedy I Inaugural Meeting of the International Computing Advisory Committee

The Fermilab International Computing Advisory Committee (ICAC) reviews and advises the laboratory on computing operations, cyber security, upgrade plans, and software and computing R&D aimed towards the development and exploitation of future facilities as well as advancing scientific tools and methods in general. The committee meets at least annually and reports to the Fermilab Director. The Fermilab Director provides a charge to the ICAC in advance of each meeting. An agenda for each meeting is negotiated between the ICAC chair and the Fermilab Chief Information Officer. The ICAC is expected to review and to monitor progress within the Fermilab computing operations, cyber security, upgrades, projects and R&D programs with respect to the established laboratory objectives, currently encompassing: Software and Computing for the Intensity Frontier Experiments; Fermilab's involvement in the HL-LHC Software and Computing Upgrades; Progress toward common solutions for the above domains; National and International cooperation and collaboration with partner institutions; The ICAC is expected to address high-level strategic, programmatic, and planning issues, rather than specific implementation details. The committee will be allocated time to discuss amongst themselves, in a closed session, any issue without laboratory management present. The ICAC will issue a written report containing an evaluation of progress, and recommendations where warranted, to the Fermilab Director within a month of each meeting. The report may contain information that does not address the charge directly, but the Chair considers valuable advice to the laboratory.

Inaugural Meeting of the International Computing Advisory Committee

from Thursday, March 14, 2019 at **08:00** to Friday, March 15, 2019 at **18:00** (US/Central)
at **Wilson Hall**

Description Remote Connection available
Physically on the 2nd floor of Wilson Hall

Thursday, March 14, 2019

08:00 - 08:10	Welcome from Nigel *10'* Speaker: Nigel Lockyer (Fermilab)
08:10 - 08:30	Breakfast
08:30 - 09:00	Welcome and Introduction of the Fermilab Computing Sector *30'* Speaker: Elizabeth Sexton-Kennedy (FNAL) Material: Slides
09:00 - 09:30	Introduction to the Fermilab Scientific Program *30'* Speaker: Dr. James Amundson (Fermilab) Material: Slides
09:30 - 10:00	Meeting Goals and Review of the Charge *30'* Speaker: Elizabeth Sexton-Kennedy (FNAL) Material: Slides
10:00 - 10:30	Coffee Break
10:30 - 11:00	HPC Strategy & US Exascale Program *30'* Speaker: Dr. James Amundson (Fermilab) Material: Slides
11:00 - 11:30	International Cooperation Strategy *30'* Speaker: Elizabeth Sexton-Kennedy (FNAL) Material: Slides
11:30 - 12:00	Discussion *30'*
12:00 - 13:00	Lunch
13:00 - 13:30	Cyber Security and Other DOE Mandates *30'* Speaker: Jon Bakken (Fermilab)