Recent Progress on Application Development in ECP

Andrew Siegel, Director of ECP Application Development

Erik Draeger, Deputy Director of ECP Application Development

Big Picture

Where we started

National security	Energy security	Economic security	Scientific discovery	Earth system	Health care
Next-generation, **stockpile stewardship** codes **Reentry-vehicle-environment** simulation Multi-physics science simulations **of high-energy density physics** conditions	Turbine **wind plant** efficiency Design and commercialization Nuclear fission and fusion reactor **materials design** Subsurface use for carbon capture, petroleum extraction, waste disposal High-efficiency, low-emission **combustion engine** and gas turbine design Scale up of **clean fossil fuel** combustion **Biofuel** catalyst design	**Additive manufacturing** of qualifiable metal parts Reliable and efficient planning of the **power grid** **Seismic** hazard risk assessment	**Cosmological probe** of the standard model of particle physics Validate fundamental laws of nature **Plasma wakefield accelerator** design Light source-enabled **analysis of protein and molecular structure** and design Find, predict, and control materials and properties Predict and control **magnetically confined fusion plasmas** Demystify **origin of chemical elements**	Accurate regional impact assessments in **Earth system models** Stress-resistant crop analysis and catalytic conversion of **biomass-derived alcohols** **Metagenomics** for analysis of biogeochemical cycles, climate change, environmental remediation	Accelerate and translate **cancer research** (partnership with NIH)

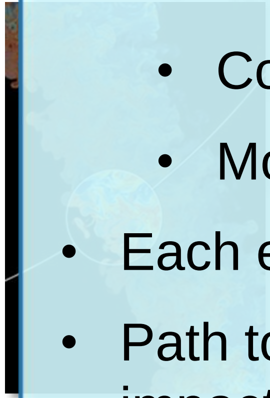

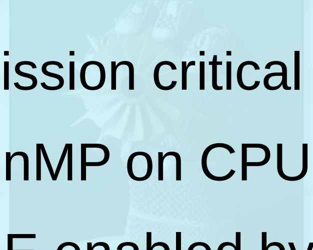

- **25 applications** and **6 co-design** projects
  - Including **51 separate codes**
  - Representing over **10 million lines of code**
  - Many supporting large user communities
  - Covering broad range of mission critical S&E domains
  - Mostly all MPI or MPI+OpenMP on CPUs
- Each envisioned innovative S&E enabled by 100X increase in computing power
- Path to harnessing 100-fold improvement initially unknown likely to have disruptive impact on software unlike anything in last 30 years

→ **Massive software investments**

EXASCALE COMPUTING PROJECT

Where we are now

- Significant progress on multi-GPU nodes across all project, particularly on **Summit** and **Sierra,** speedups from 7-200X baseline
- **Co-design Centers** have surpassed original vision, developed into best practice
- Refactoring code for heterogeneous machine has required **fundamental changes to data structures**, **data movement and algorithms** that independent of specific accelerator features.
- AD projects are guinea pigs in exercising **performance portable programming** models

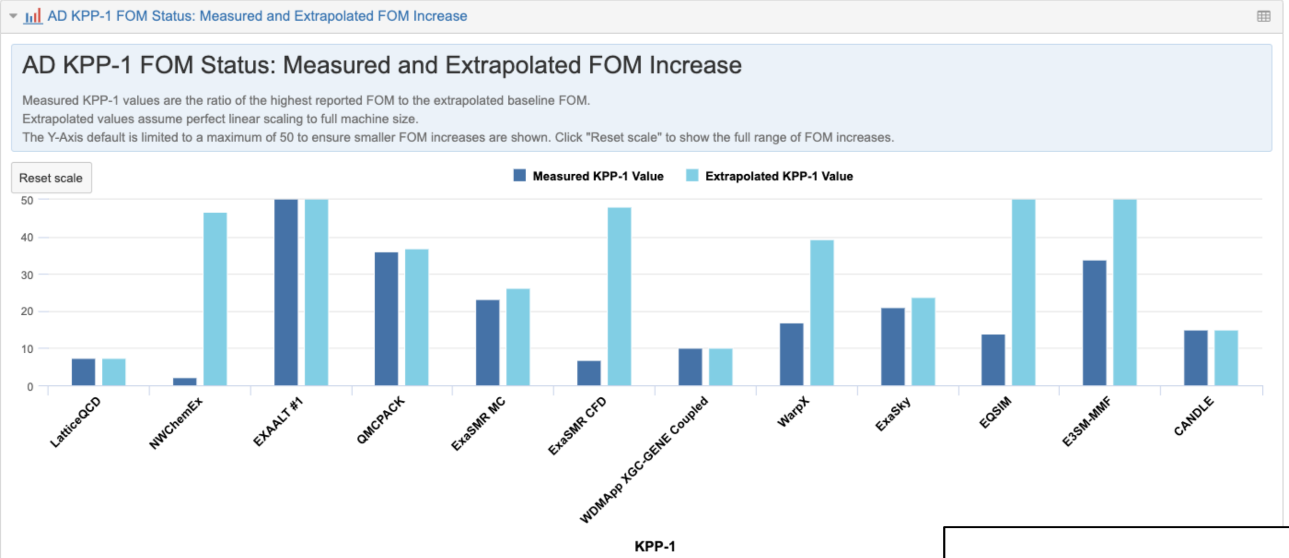
AD KPP-1 FOM Status: Measured and Extrapolated FOM Increase

AD KPP-1 FOM Status: Measured and Extrapolated FOM Increase

Measured KPP-1 values are the ratio of the highest reported FOM to the extrapolated baseline FOM.
Extrapolated values assume perfect linear scaling to full machine size.
The Y-Axis default is limited to a maximum of 50 to ensure smaller FOM increases are shown. Click "Reset scale" to show the full range of FOM increases.


ECP-U-2017-XXX

State and Progress of ECP Applications: Application Development

WBS 2.2, Milestone ECP-XX-XXXX

Andrew Siegel[1], Erik Draeger[2], Jack Deslippe[3], Anshu Dubey[1], Tom Evans[4], Tim Germann[5], and William Hart[6]

[1]Argonne National Laboratory
[2]Lawrence Livermore National Laboratory
[3]Lawrence Berkeley National Laboratory
[4]Oak Ridge National Laboratory
[5]Los Alamos National Laboratory
[6]Sandia National Laboratories

January 14, 2020

FY19 ECP AD Assessment Report

Where we are going

Department of Energy (DOE) Roadmap to Exascale Systems

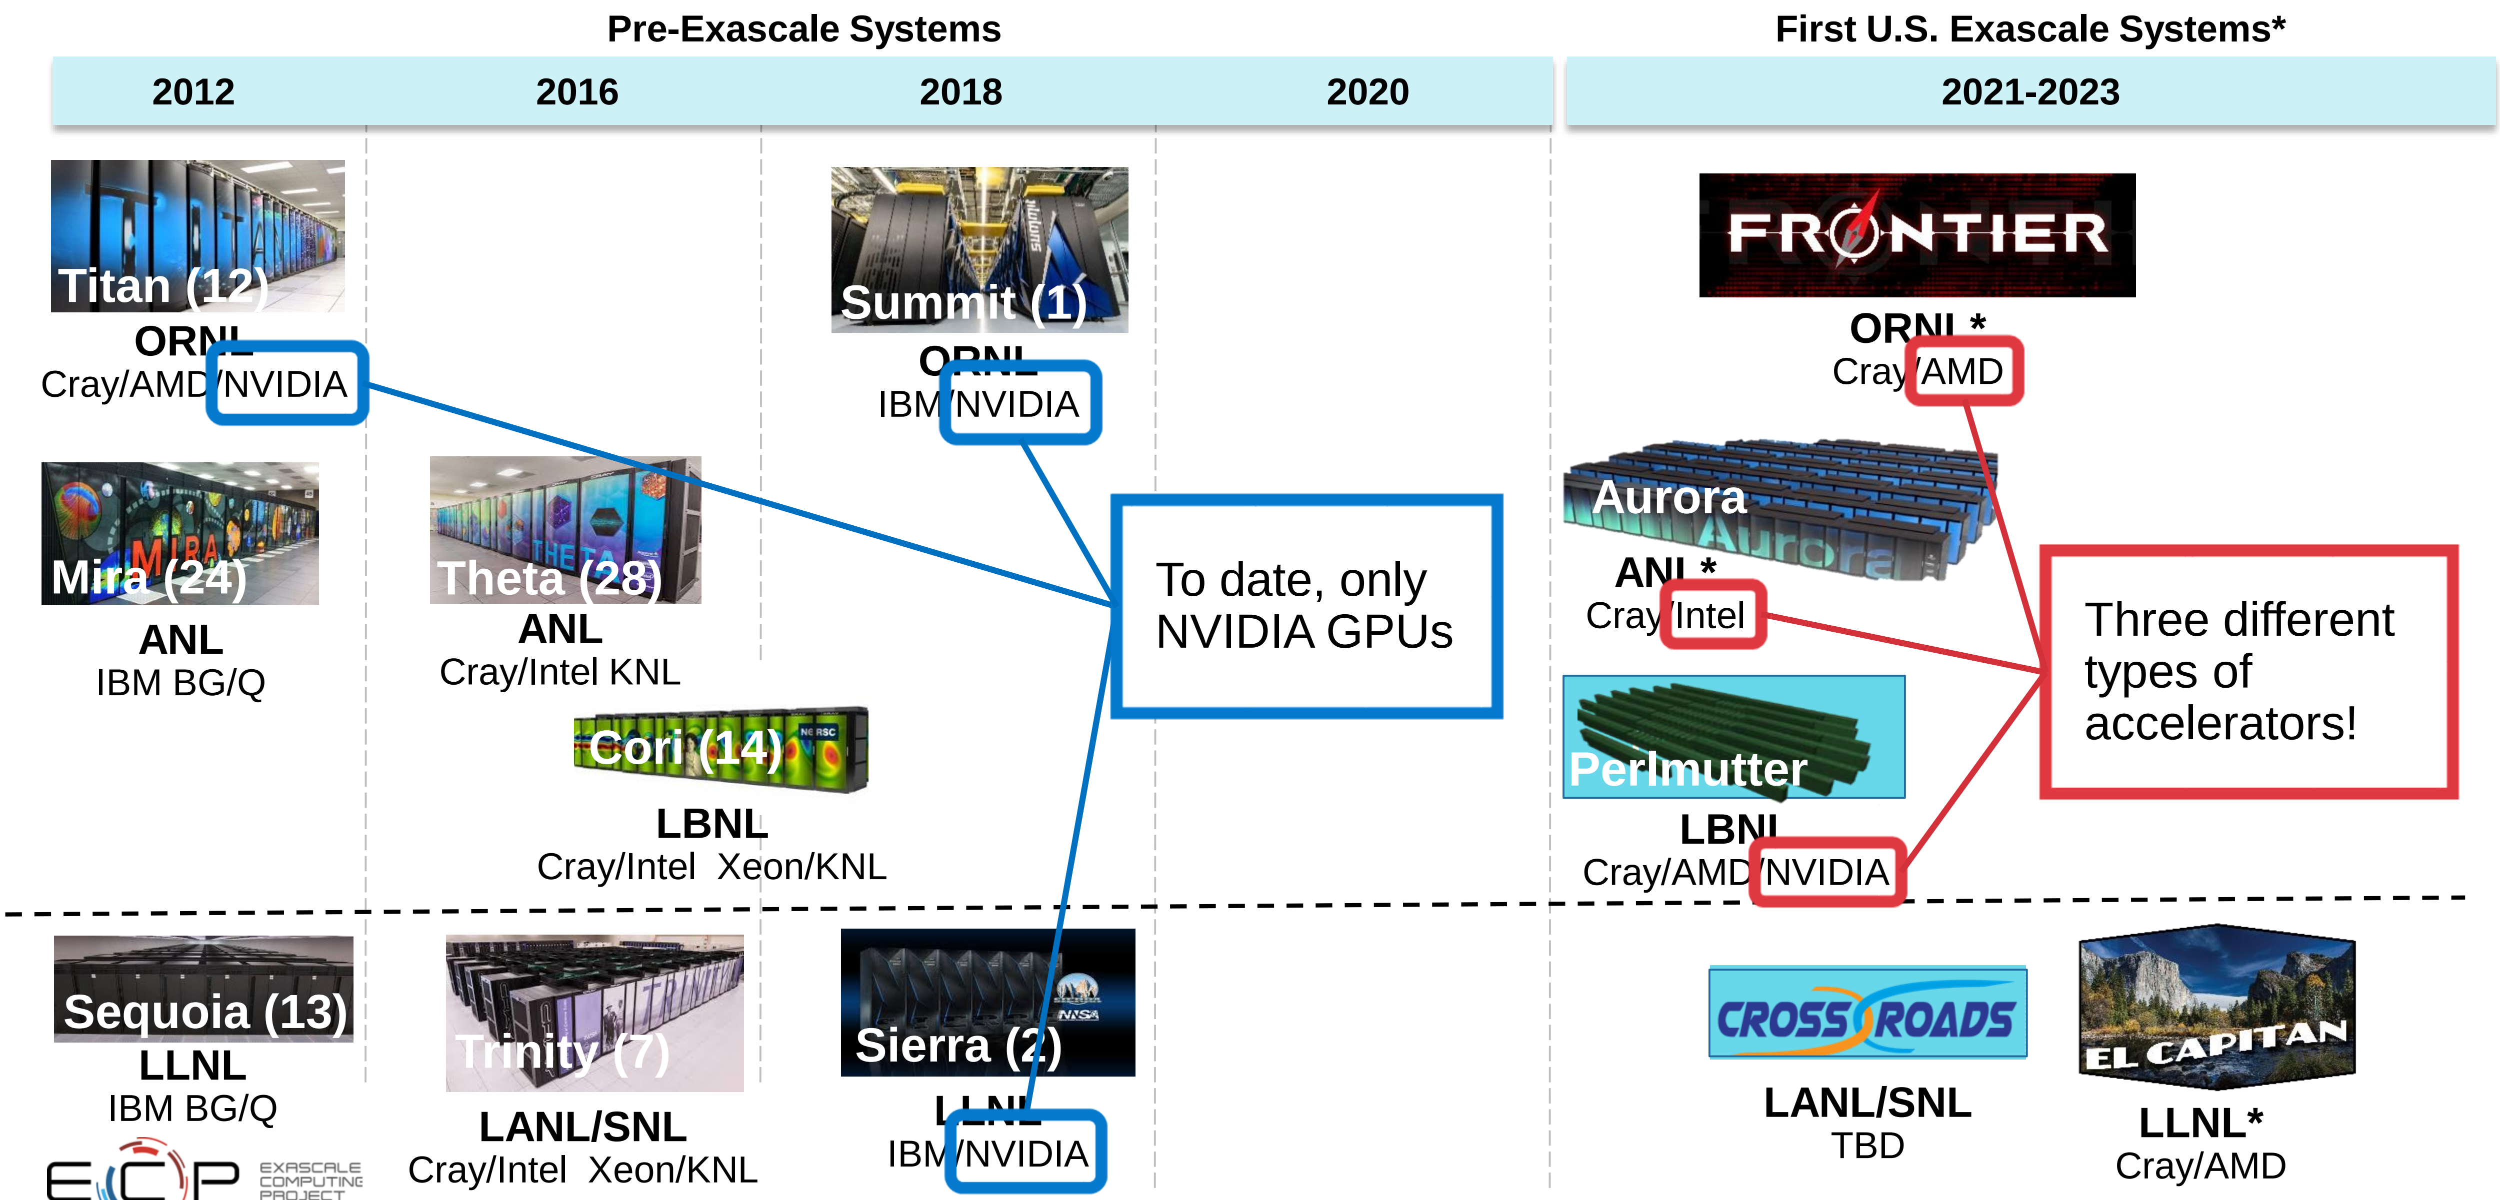

Pre-Exascale Systems

2012 | **2016** | **2018** | **2020**

First U.S. Exascale Systems*

2021-2023

ORNL
Cray/AMD/NVIDIA

ANL
IBM BG/Q

ANL
Cray/Intel KNL

LBNL
Cray/Intel Xeon/KNL

ORNL
IBM/NVIDIA

LLNL
IBM BG/Q

LANL/SNL
Cray/Intel Xeon/KNL

LLNL
IBM/NVIDIA

ORNL *
Cray/AMD

ANL *
Cray/Intel

LBNL
Cray/AMD/NVIDIA

LANL/SNL
TBD

LLNL*
Cray/AMD

To date, only NVIDIA GPUs

Three different types of accelerators!

Early access hardware

Tulip

Frontier Center of Excellence System

- 8 Compute nodes, each with:
  - 1x AMD EPYC 7601(32C/180W/2.2GHZ)
  - 256GB 2666 DDR Memory
  - 1x ConnectX-5 EDR adapter
  - 1x 480GB SSD
- 6 of the nodes have AMD GPUs:
  - 4x AMD MI60 32GB 300W GPU PCIe
- 2 of the nodes have Nvidia GPUs:
  - 4x NVIDIA V100 32GB 250W GPU PCIe

Iris

Aurora Center of Excellence System

- 20 Compute nodes, each with:
  - 1x Intel Xeon E3-1585 v5 CPU w/ Intel Iris Pro Graphics P580 (Intel Gen9 GPU)
  - 64GB DDR4 (operating at DDR4-2133)
  - 1Gbit ethernet
  - OneAPI beta SDK
  - /home, /soft NFS mounted storage

Bird’s-eye View Application Development Timeline

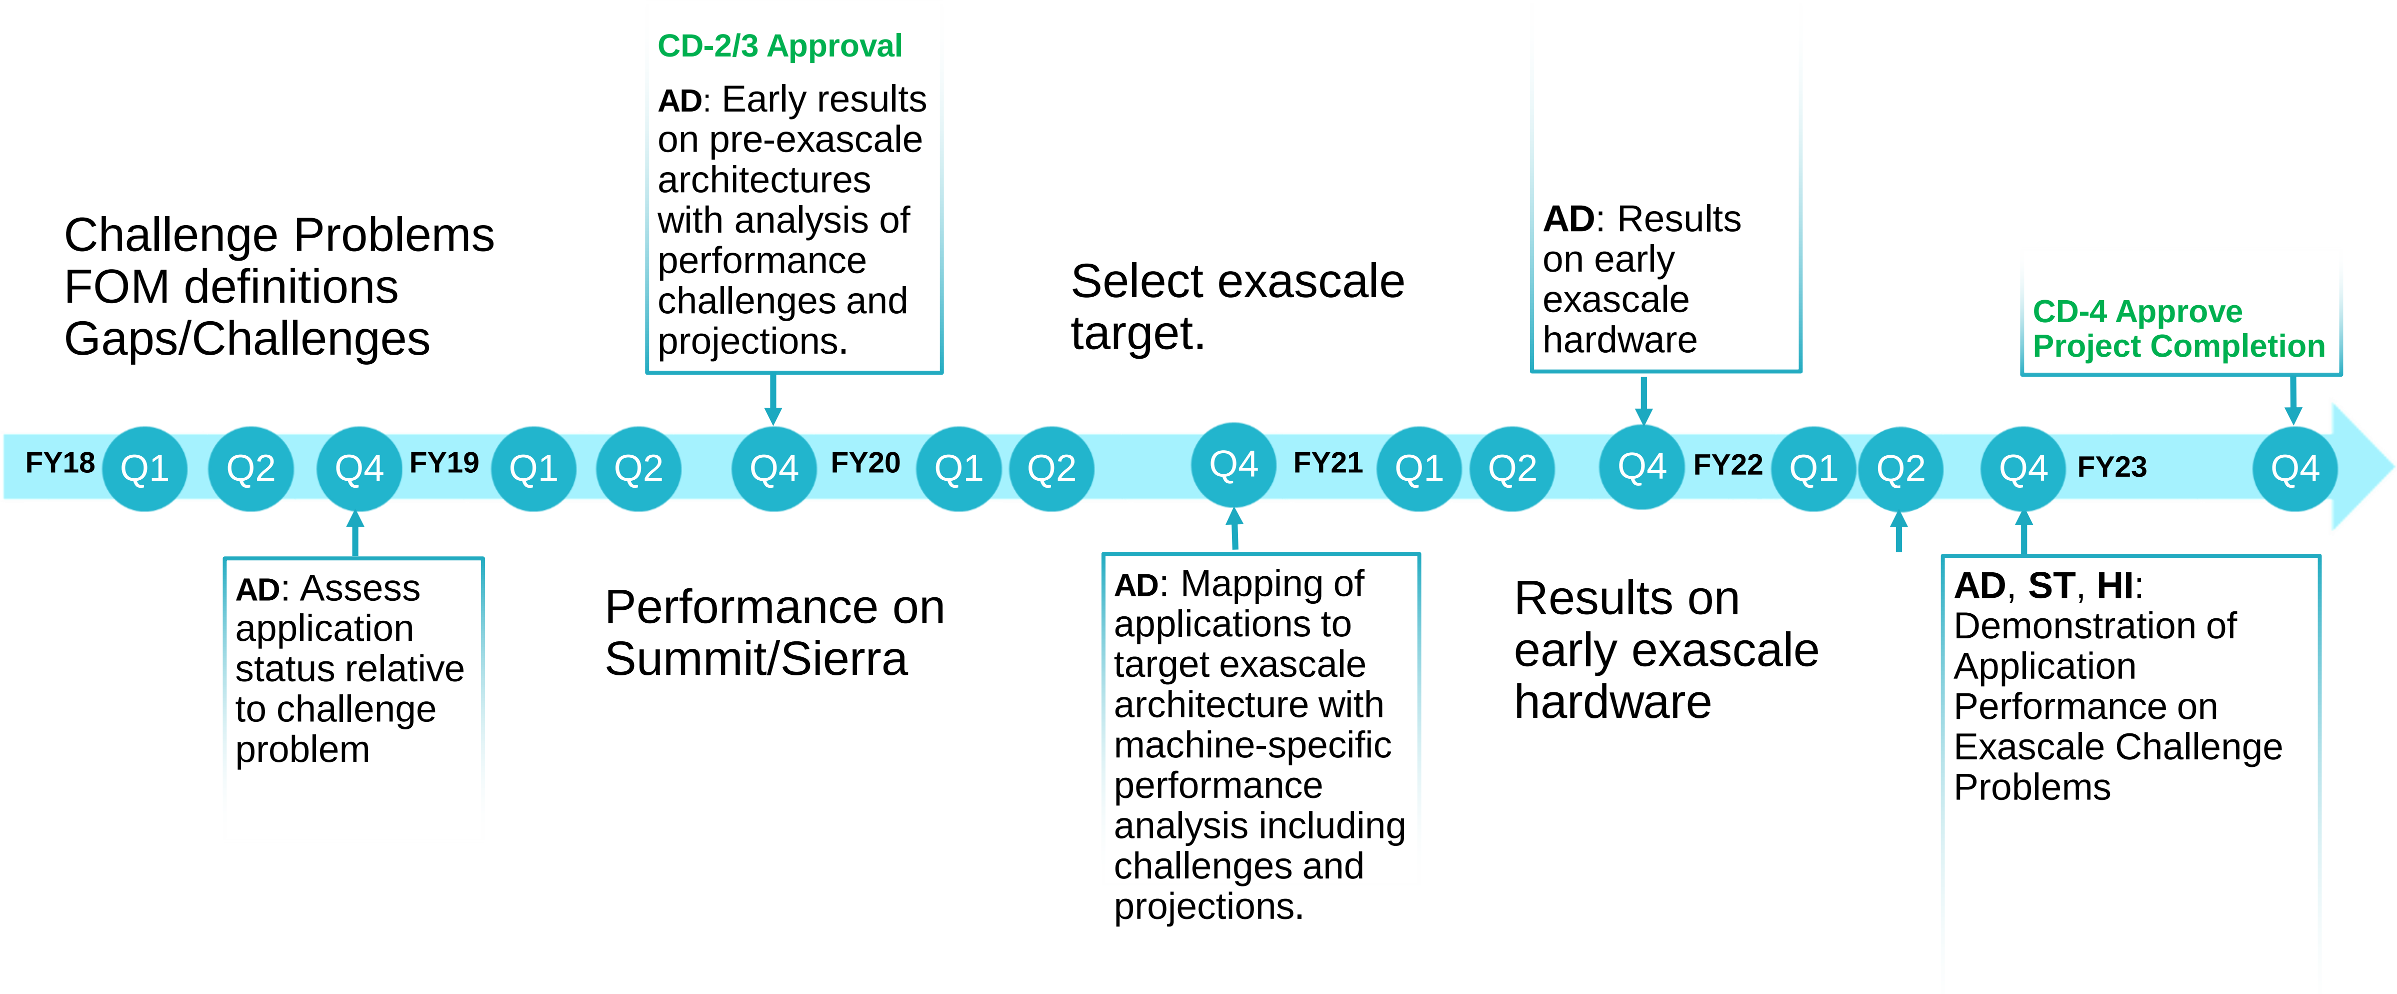

Challenge Problems
FOM definitions
Gaps/Challenges

AD: Assess application status relative to challenge problem

CD-2/3 Approval

AD: Early results on pre-exascale architectures with analysis of performance challenges and projections.

Performance on Summit/Sierra

Select exascale target.

AD: Mapping of applications to target exascale architecture with machine-specific performance analysis including challenges and projections.

AD: Results on early exascale hardware

Results on early exascale hardware

AD, **ST**, **HI**: Demonstration of Application Performance on Exascale Challenge Problems

CD-4 Approve Project Completion

We have committed to quantified Key Performance Parameters (KPPs)

KPP ID	Description of Scope	Threshold KPP	Objective KPP	Verification Action/Evidence
KPP-1	Performance improvement for mission-critical problems	50% of selected applications achieve Figure of Merit improvement ≥50	100% of selected applications achieve their KPP-1 stretch goal	Independent assessment of measured results and report that threshold goal is met
KPP-2	Broaden the reach of exascale science and mission capability	50% of selected applications can execute their challenge problem	100% of selected applications can execute their challenge problem stretch goal	Independent assessment of mission application readiness
KPP-3	Productive and sustainable software ecosystem	50% of the weighted impact goals are met	100% of the weighted impact stretch goals are met	Independent assessment verifying threshold goal is met
KPP-4	Enrich the HPC hardware ecosystem	Vendors meet 80% of all the PathForward milestones	Vendors meet 100% of all the PathForward milestones	Independent review of the PathForward milestones to assure they meet the contract requirements; evidence is the final milestone deliverable

11 apps

13 apps

70 S/W products

AMD
Cray
HPE
IBM
Intel
NVIDIA

EXASCALE COMPUTING PROJECT

Measuring Progress: KPP-1

Figure of Merit (FOM) Dashboard

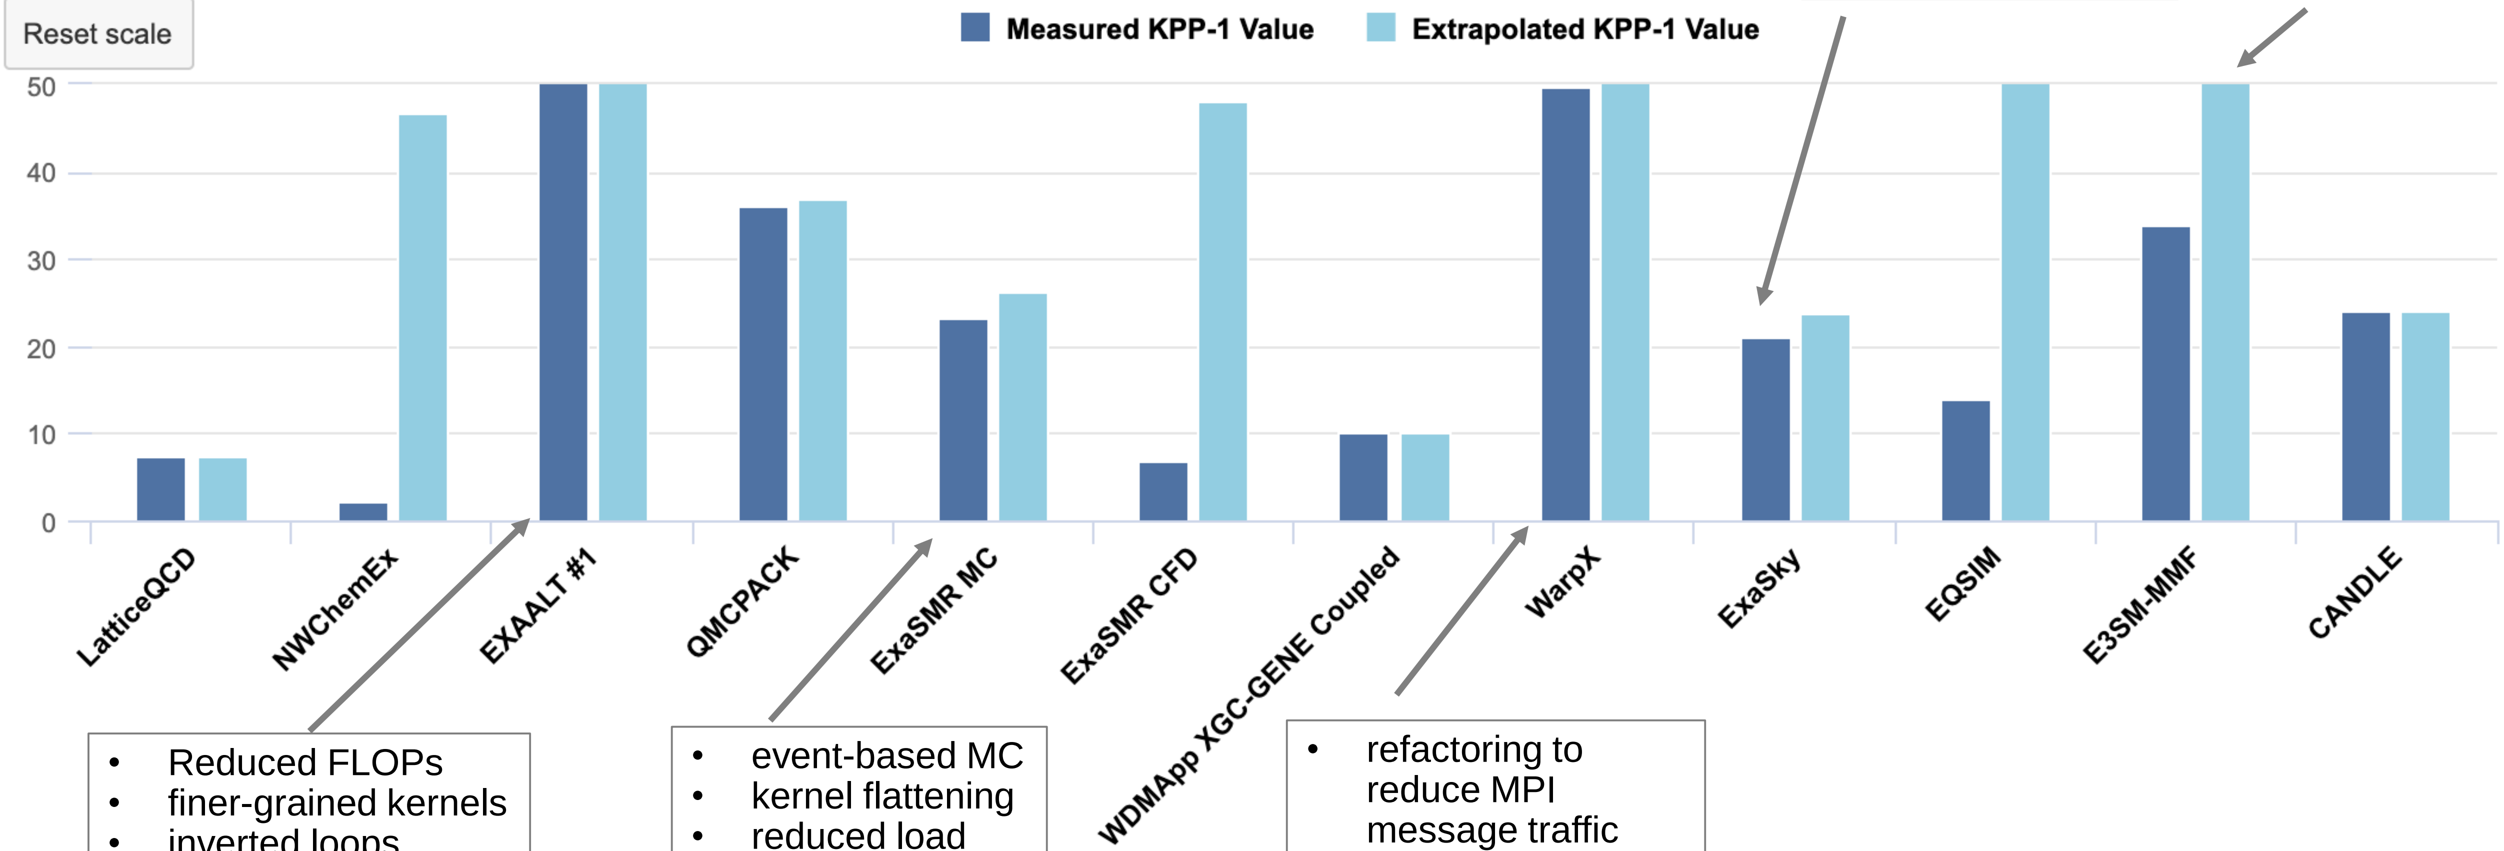

GPU-resident
Version of first-order
CRK SPH.

Switch to MMF
Use of GPUs

- Reduced FLOPs
- finer-grained kernels
- inverted loops
- flattened 3d arrays
- Ported to Kokkos

- event-based MC
- kernel flattening
- reduced load imbalances among threads

- refactoring to reduce MPI message traffic
- Increased Summit node count

ExaSMR FOM updates

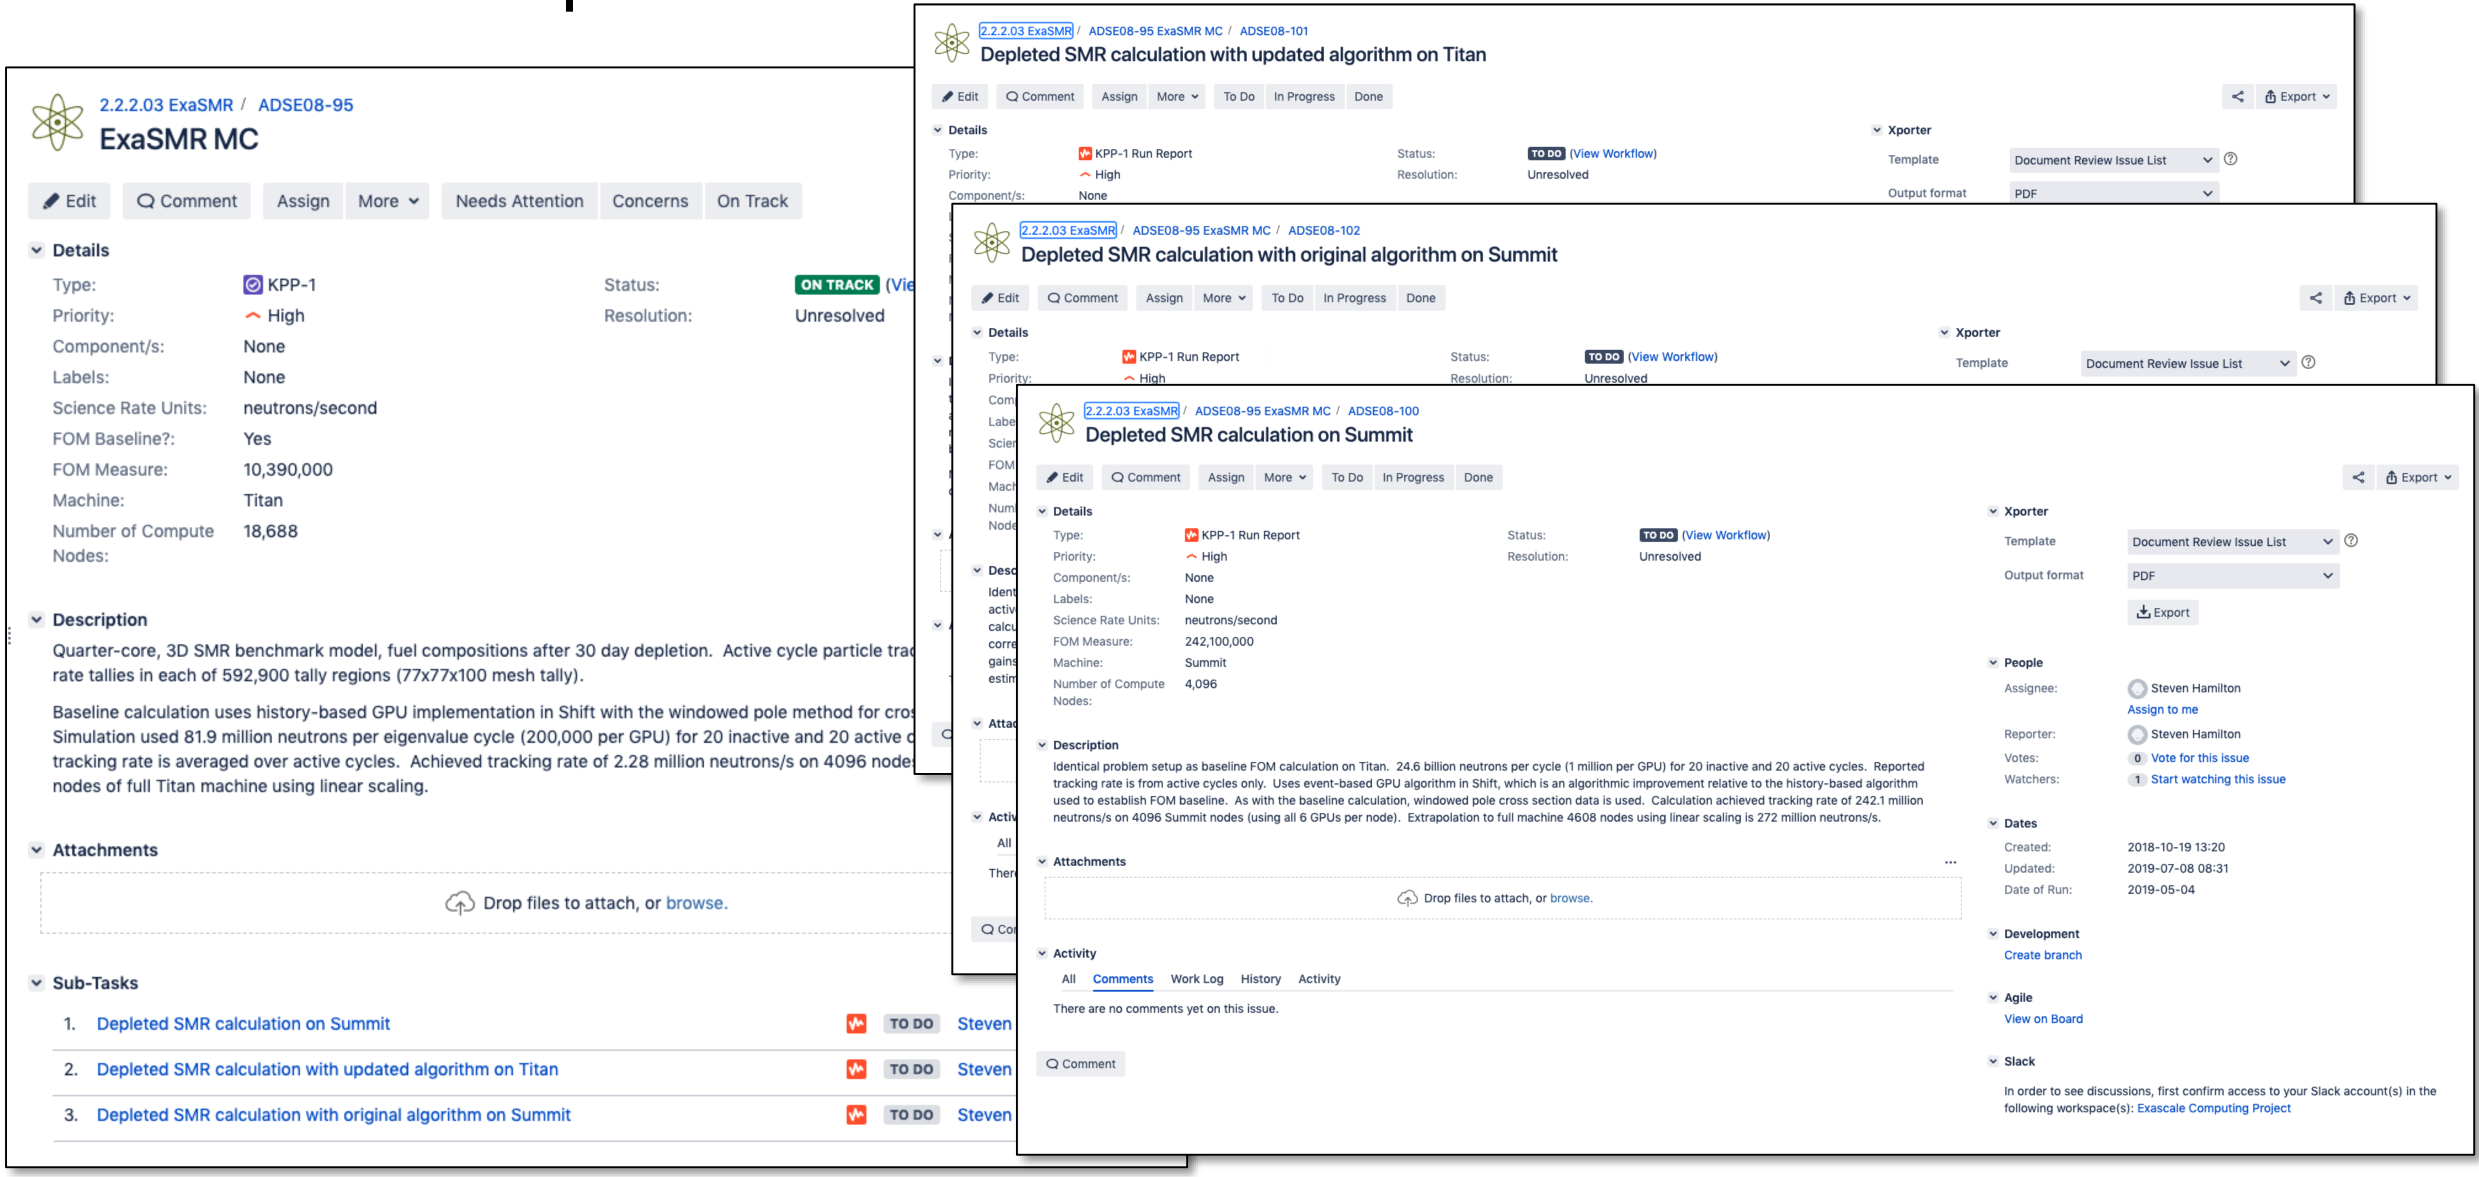

2.2.2.03 ExaSMR / ADSE08-95

ExaSMR MC

Details

Type:	KPP-1	Status:	ON TRACK (Vie
Priority:	High	Resolution:	Unresolved
Component/s:	None		
Labels:	None		
Science Rate Units:	neutrons/second		
FOM Baseline?:	Yes		
FOM Measure:	10,390,000		
Machine:	Titan		
Number of Compute Nodes:	18,688		

Description

Quarter-core, 3D SMR benchmark model, fuel compositions after 30 day depletion. Active cycle particle trac
rate tallies in each of 592,900 tally regions (77x77x100 mesh tally).

Baseline calculation uses history-based GPU implementation in Shift with the windowed pole method for cros
Simulation used 81.9 million neutrons per eigenvalue cycle (200,000 per GPU) for 20 inactive and 20 active c
tracking rate is averaged over active cycles. Achieved tracking rate of 2.28 million neutrons/s on 4096 node
nodes of full Titan machine using linear scaling.

Attachments

Drop files to attach, or browse.

Sub-Tasks

1. Depleted SMR calculation on Summit — TO DO — Steven
2. Depleted SMR calculation with updated algorithm on Titan — TO DO — Steven
3. Depleted SMR calculation with original algorithm on Summit — TO DO — Steven

2.2.2.03 ExaSMR / ADSE08-95 ExaSMR MC / ADSE08-101

Depleted SMR calculation with updated algorithm on Titan

Details

Type:	KPP-1 Run Report	Status:	TO DO (View Workflow)
Priority:	High	Resolution:	Unresolved
Component/s:	None		

Xporter

Template: Document Review Issue List

Output format: PDF

2.2.2.03 ExaSMR / ADSE08-95 ExaSMR MC / ADSE08-102

Depleted SMR calculation with original algorithm on Summit

Details

Type:	KPP-1 Run Report	Status:	TO DO (View Workflow)
Priority:	High	Resolution:	Unresolved

Xporter

Template: Document Review Issue List

2.2.2.03 ExaSMR / ADSE08-95 ExaSMR MC / ADSE08-100

Depleted SMR calculation on Summit

Details

Type:	KPP-1 Run Report	Status:	TO DO (View Workflow)
Priority:	High	Resolution:	Unresolved
Component/s:	None		
Labels:	None		
Science Rate Units:	neutrons/second		
FOM Measure:	242,100,000		
Machine:	Summit		
Number of Compute Nodes:	4,096		

Description

Identical problem setup as baseline FOM calculation on Titan. 24.6 billion neutrons per cycle (1 million per GPU) for 20 inactive and 20 active cycles. Reported tracking rate is from active cycles only. Uses event-based GPU algorithm in Shift, which is an algorithmic improvement relative to the history-based algorithm used to establish FOM baseline. As with the baseline calculation, windowed pole cross section data is used. Calculation achieved tracking rate of 242.1 million neutrons/s on 4096 Summit nodes (using all 6 GPUs per node). Extrapolation to full machine 4608 nodes using linear scaling is 272 million neutrons/s.

Attachments

Drop files to attach, or browse.

Activity

There are no comments yet on this issue.

Xporter

Template: Document Review Issue List

Output format: PDF

People

Assignee: Steven Hamilton

Assign to me

Reporter: Steven Hamilton

Votes: 0 Vote for this issue

Watchers: 1 Start watching this issue

Dates

Created: 2018-10-19 13:20

Updated: 2019-07-08 08:31

Date of Run: 2019-05-04

Development

Create branch

Agile

View on Board

Slack

In order to see discussions, first confirm access to your Slack account(s) in the following workspace(s): Exascale Computing Project

EXASCALE COMPUTING PROJECT

Example of JIRA issues entered by PI to log FOM calculation

WarpX FOM updates

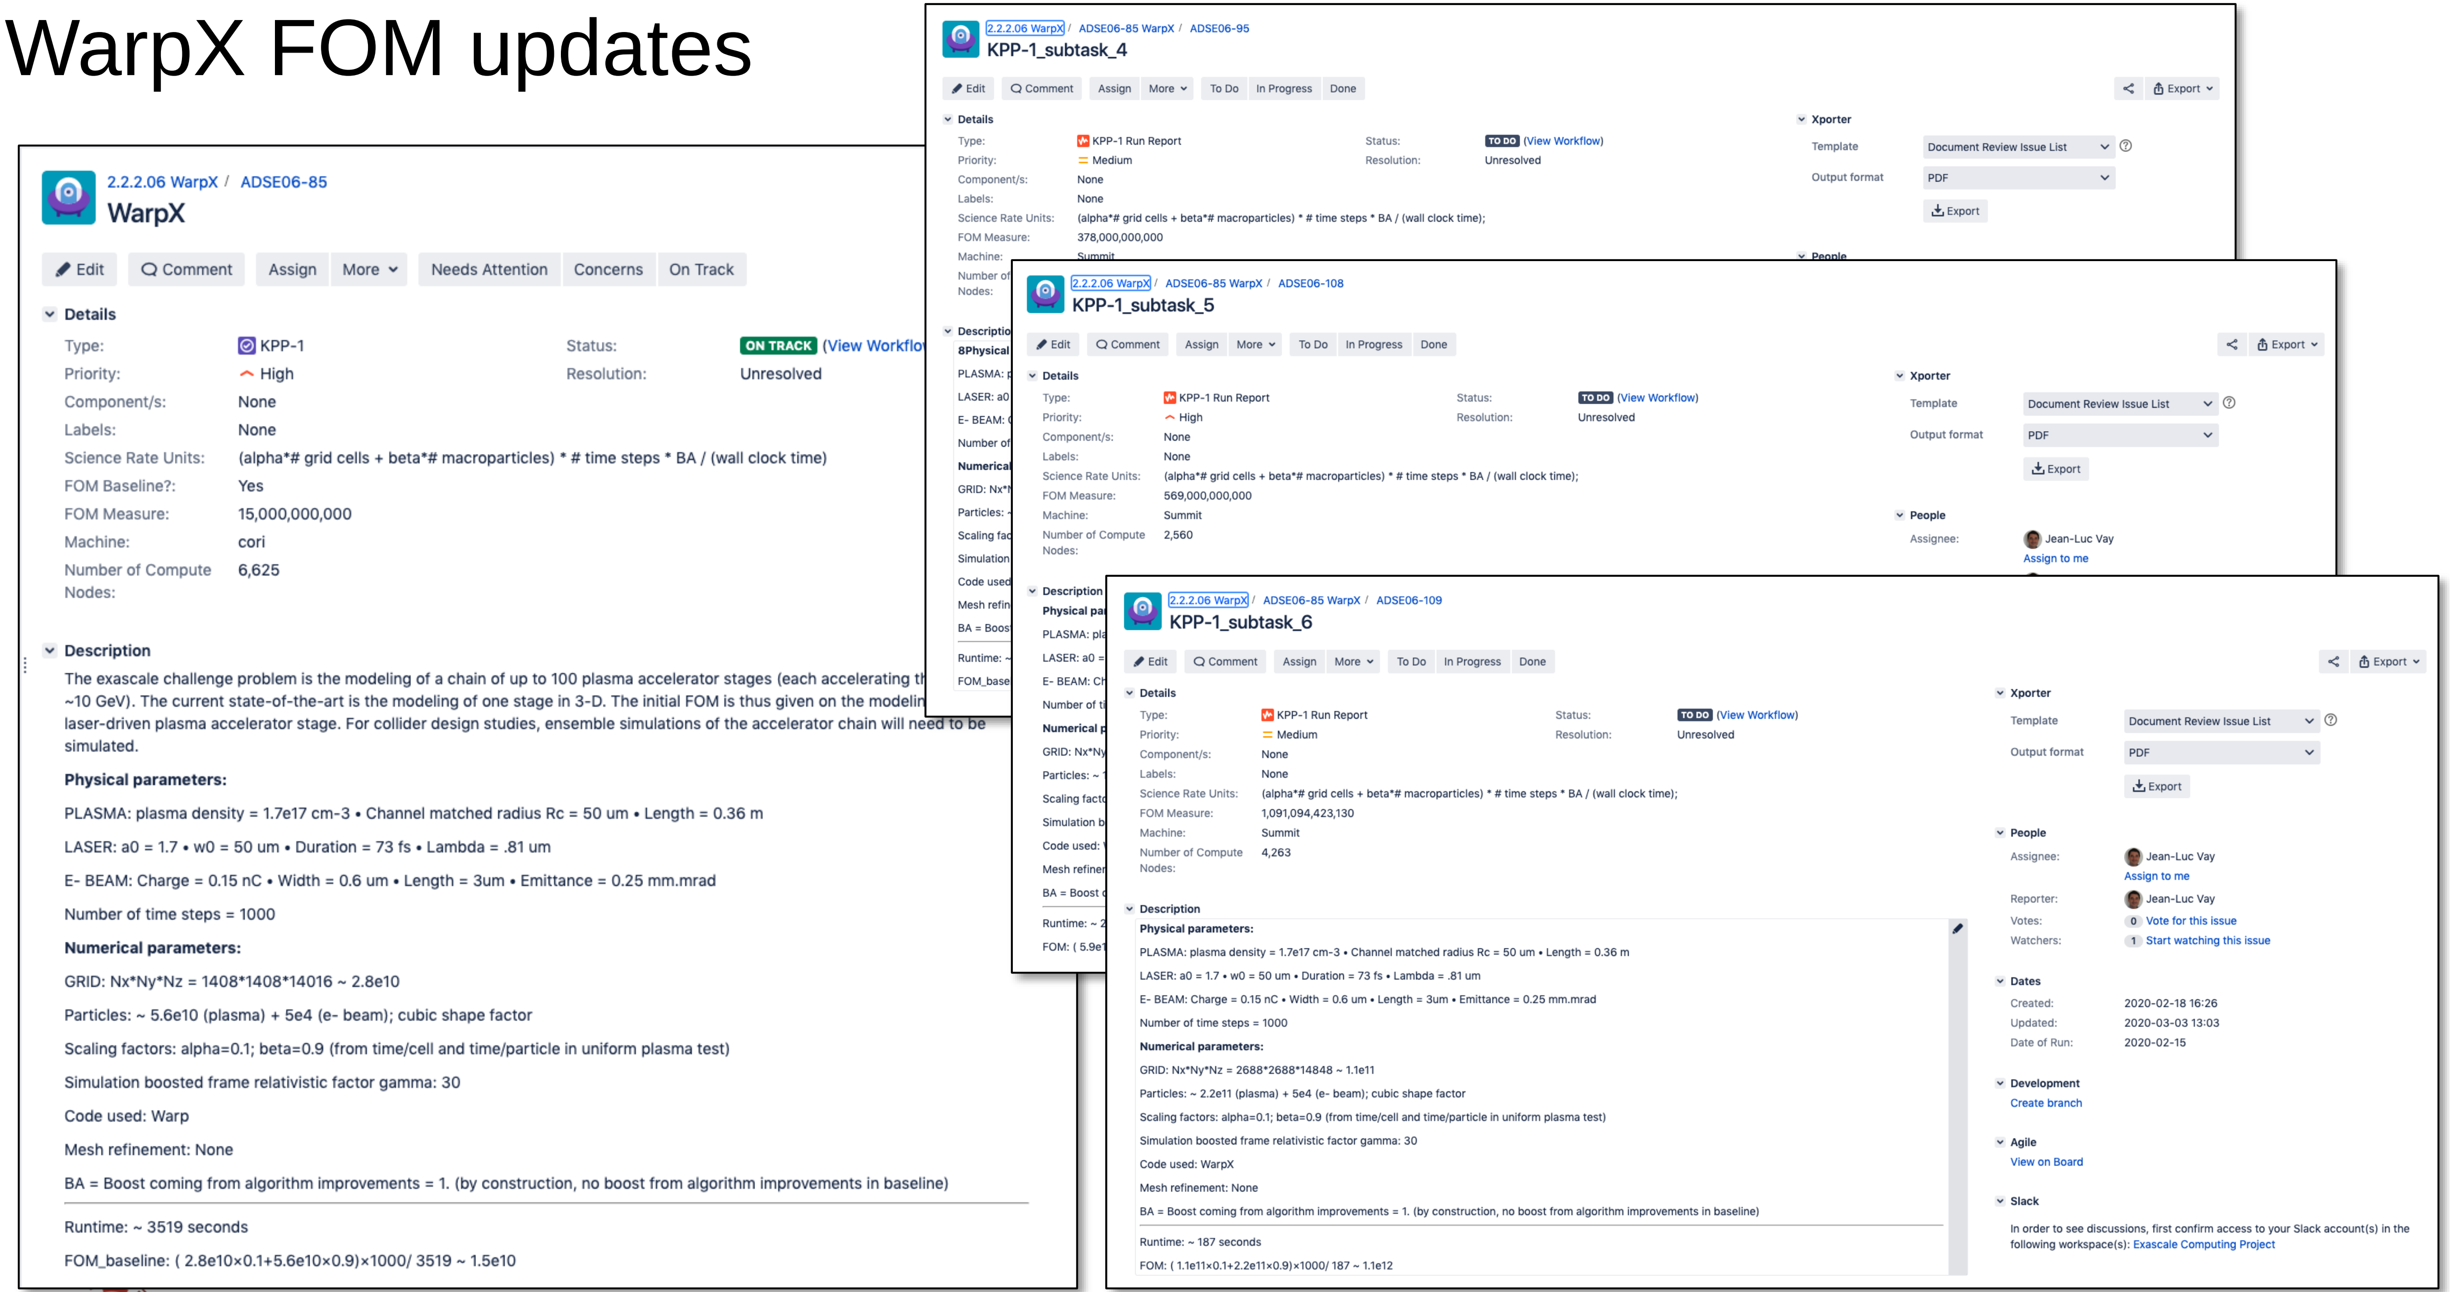

WarpX

Details

Type: KPP-1
Priority: High
Component/s: None
Labels: None
Science Rate Units: (alpha*# grid cells + beta*# macroparticles) * # time steps * BA / (wall clock time)
FOM Baseline?: Yes
FOM Measure: 15,000,000,000
Machine: cori
Number of Compute Nodes: 6,625
Status: ON TRACK (View Workflo
Resolution: Unresolved

Description

The exascale challenge problem is the modeling of a chain of up to 100 plasma accelerator stages (each accelerating t ~10 GeV). The current state-of-the-art is the modeling of one stage in 3-D. The initial FOM is thus given on the modelin laser-driven plasma accelerator stage. For collider design studies, ensemble simulations of the accelerator chain will need to be simulated.

Physical parameters:

PLASMA: plasma density = 1.7e17 cm-3 • Channel matched radius Rc = 50 um • Length = 0.36 m

LASER: a0 = 1.7 • w0 = 50 um • Duration = 73 fs • Lambda = .81 um

E- BEAM: Charge = 0.15 nC • Width = 0.6 um • Length = 3um • Emittance = 0.25 mm.mrad

Number of time steps = 1000

Numerical parameters:

GRID: Nx*Ny*Nz = 1408*1408*14016 ~ 2.8e10

Particles: ~ 5.6e10 (plasma) + 5e4 (e- beam); cubic shape factor

Scaling factors: alpha=0.1; beta=0.9 (from time/cell and time/particle in uniform plasma test)

Simulation boosted frame relativistic factor gamma: 30

Code used: Warp

Mesh refinement: None

BA = Boost coming from algorithm improvements = 1. (by construction, no boost from algorithm improvements in baseline)

Runtime: ~ 3519 seconds

FOM_baseline: (2.8e10×0.1+5.6e10×0.9)×1000/ 3519 ~ 1.5e10

KPP-1_subtask_4

Details

Type: KPP-1 Run Report
Priority: Medium
Component/s: None
Labels: None
Science Rate Units: (alpha*# grid cells + beta*# macroparticles) * # time steps * BA / (wall clock time);
FOM Measure: 378,000,000,000
Machine: Summit
Status: TO DO (View Workflow)
Resolution: Unresolved

KPP-1_subtask_5

Details

Type: KPP-1 Run Report
Priority: High
Component/s: None
Labels: None
Science Rate Units: (alpha*# grid cells + beta*# macroparticles) * # time steps * BA / (wall clock time);
FOM Measure: 569,000,000,000
Machine: Summit
Number of Compute Nodes: 2,560
Status: TO DO (View Workflow)
Resolution: Unresolved
Assignee: Jean-Luc Vay

KPP-1_subtask_6

Details

Type: KPP-1 Run Report
Priority: Medium
Component/s: None
Labels: None
Science Rate Units: (alpha*# grid cells + beta*# macroparticles) * # time steps * BA / (wall clock time);
FOM Measure: 1,091,094,423,130
Machine: Summit
Number of Compute Nodes: 4,263
Status: TO DO (View Workflow)
Resolution: Unresolved

Assignee: Jean-Luc Vay
Reporter: Jean-Luc Vay
Votes: 0
Watchers: 1

Created: 2020-02-18 16:26
Updated: 2020-03-03 13:03
Date of Run: 2020-02-15

In order to see discussions, first confirm access to your Slack account(s) in the following workspace(s): Exascale Computing Project

Description

Physical parameters:

PLASMA: plasma density = 1.7e17 cm-3 • Channel matched radius Rc = 50 um • Length = 0.36 m

LASER: a0 = 1.7 • w0 = 50 um • Duration = 73 fs • Lambda = .81 um

E- BEAM: Charge = 0.15 nC • Width = 0.6 um • Length = 3um • Emittance = 0.25 mm.mrad

Number of time steps = 1000

Numerical parameters:

GRID: Nx*Ny*Nz = 2688*2688*14848 ~ 1.1e11

Particles: ~ 2.2e11 (plasma) + 5e4 (e- beam); cubic shape factor

Scaling factors: alpha=0.1; beta=0.9 (from time/cell and time/particle in uniform plasma test)

Simulation boosted frame relativistic factor gamma: 30

Code used: WarpX

Mesh refinement: None

BA = Boost coming from algorithm improvements = 1. (by construction, no boost from algorithm improvements in baseline)

Runtime: ~ 187 seconds

FOM: (1.1e11×0.1+2.2e11×0.9)×1000/ 187 ~ 1.1e12

Outreach

Coordinated Publication Efforts

Special Issue Journal themes (led by Julia White)

- ✓ Co-design Centers/computational motifs
  - Contributors: AMReX, CEED, Copa, CODAR, FFT, ExaGraph, ExaLearn
  - *International Journal of High Performance Computing Applications*
  - Timeline: gather articles by end of August 2020, review by Nov., publication by end of CY20
- ✓ Coupled-application codes using accelerated systems
  - Contributors: MFIX-Exa, ExaStar, EFFIS, ExaAM
  - *International Journal of High Performance Computing Applications*
  - Timeline: gather articles by end of August 2020, review by Nov., publication by end of CY20
- – Challenges and best practices for using accelerated nodes
  - One to two issues per year, multiple years
  - Timeline: Finalize contributors by March 2020

- Phil Transactions Review article, published Jan 2020: https://doi.org/10.1098/rsta.2019.0056

ECP Industry Council Deep Dive: ANL (Virtual), March 10-11

- **When does this wave hit mid-range computing?**
  - Is it inevitable?
  - What are viable alternatives in the next several years?

55 From Industry Council Member Companies

12 Non-member Company Representatives

76 ECP and DOE Labs

3 Federal Agency (1 NASA, 2 DOE)

14 unable to categorize

160 total attendees

- **How long will multi-GPU-node systems be relevant?**
  - What is next and how do these systems evolve?
  - Should I wait and see
- **What is software cost of porting to GPU architectures?**
  - Portability across GPU vendors
  - Incremental migration to GPUs

20 companies represented

15 Industry Council Member Companies

5 Other Industry

Technical Assessment
AD Annual Report

FY19 ECP AD Assessment Report

ECP-U-2017-XXX

State and Progress of ECP Applications: Application Development

WBS 2.2, Milestone ECP-XX-XXXX

Andrew Siegel[1], Erik Draeger[2], Jack Deslippe[3],
Anshu Dubey[1], Tom Evans[4], Tim Germann[5],
and William Hart[6]

[1]Argonne National Laboratory
[2]Lawrence Livermore National Laboratory
[3]Lawrence Berkeley National Laboratory
[4]Oak Ridge National Laboratory
[5]Los Alamos National Laboratory
[6]Sandia National Laboratories

January 14, 2020

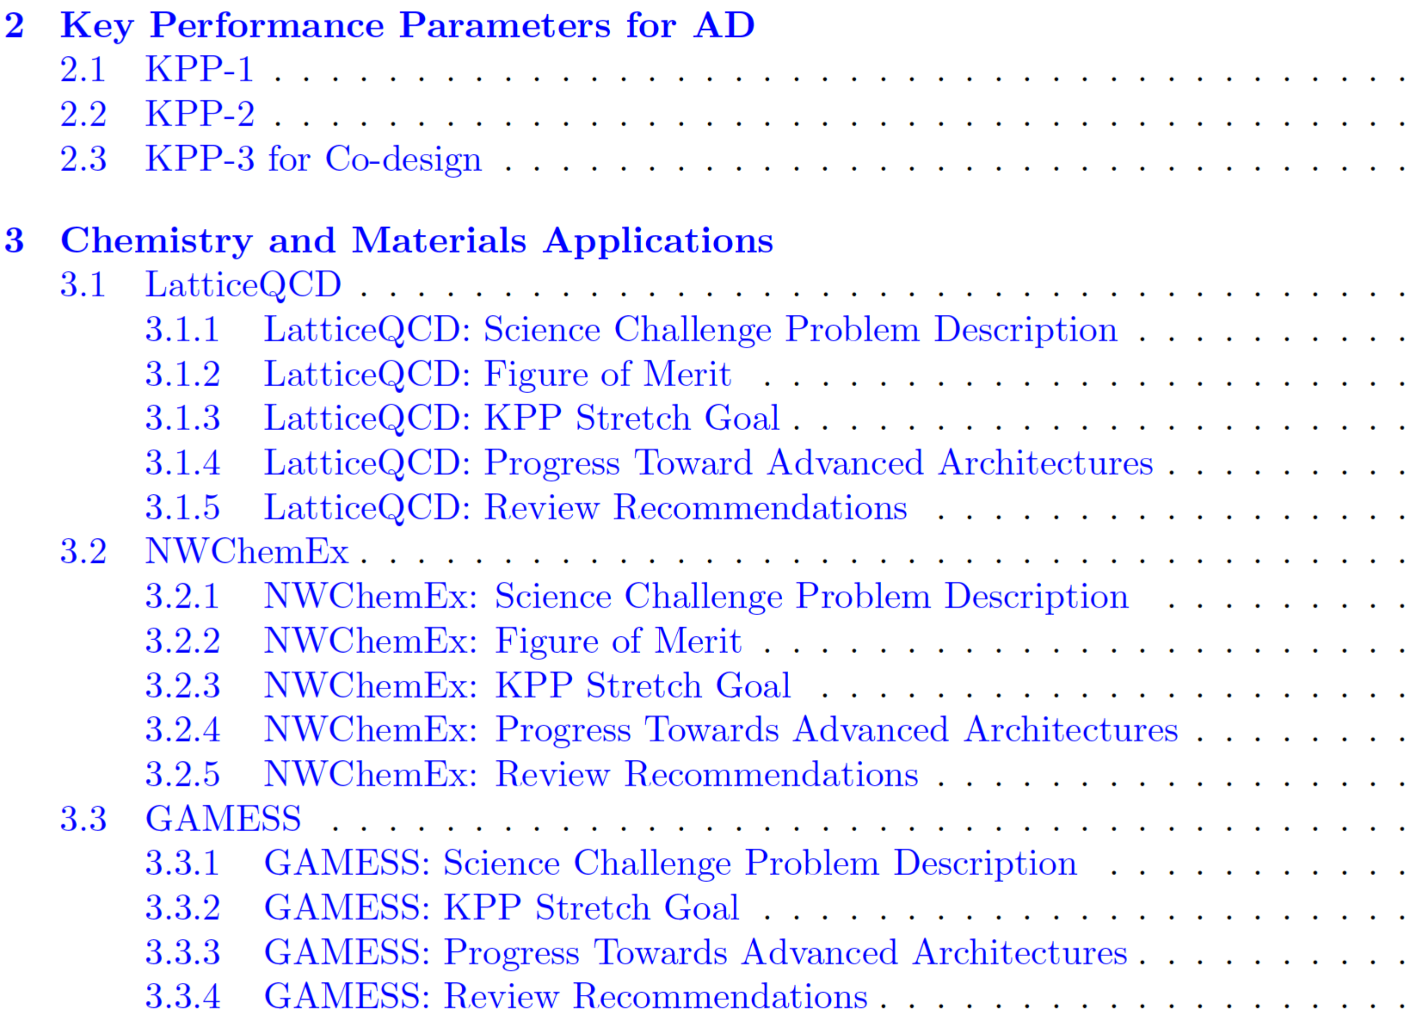

- 24 different applications
- 6 co-design projects

Common Themes Emerging from Report

1. **Flat performance profiles**
2. **Strong Scaling**
3. **Understanding/analyzing accelerator performance**
4. **Choice of programming model**
5. **Selecting mathematical models that fit architecture**
6. **Managing software dependencies**

3) Understanding/analyzing accelerator performance

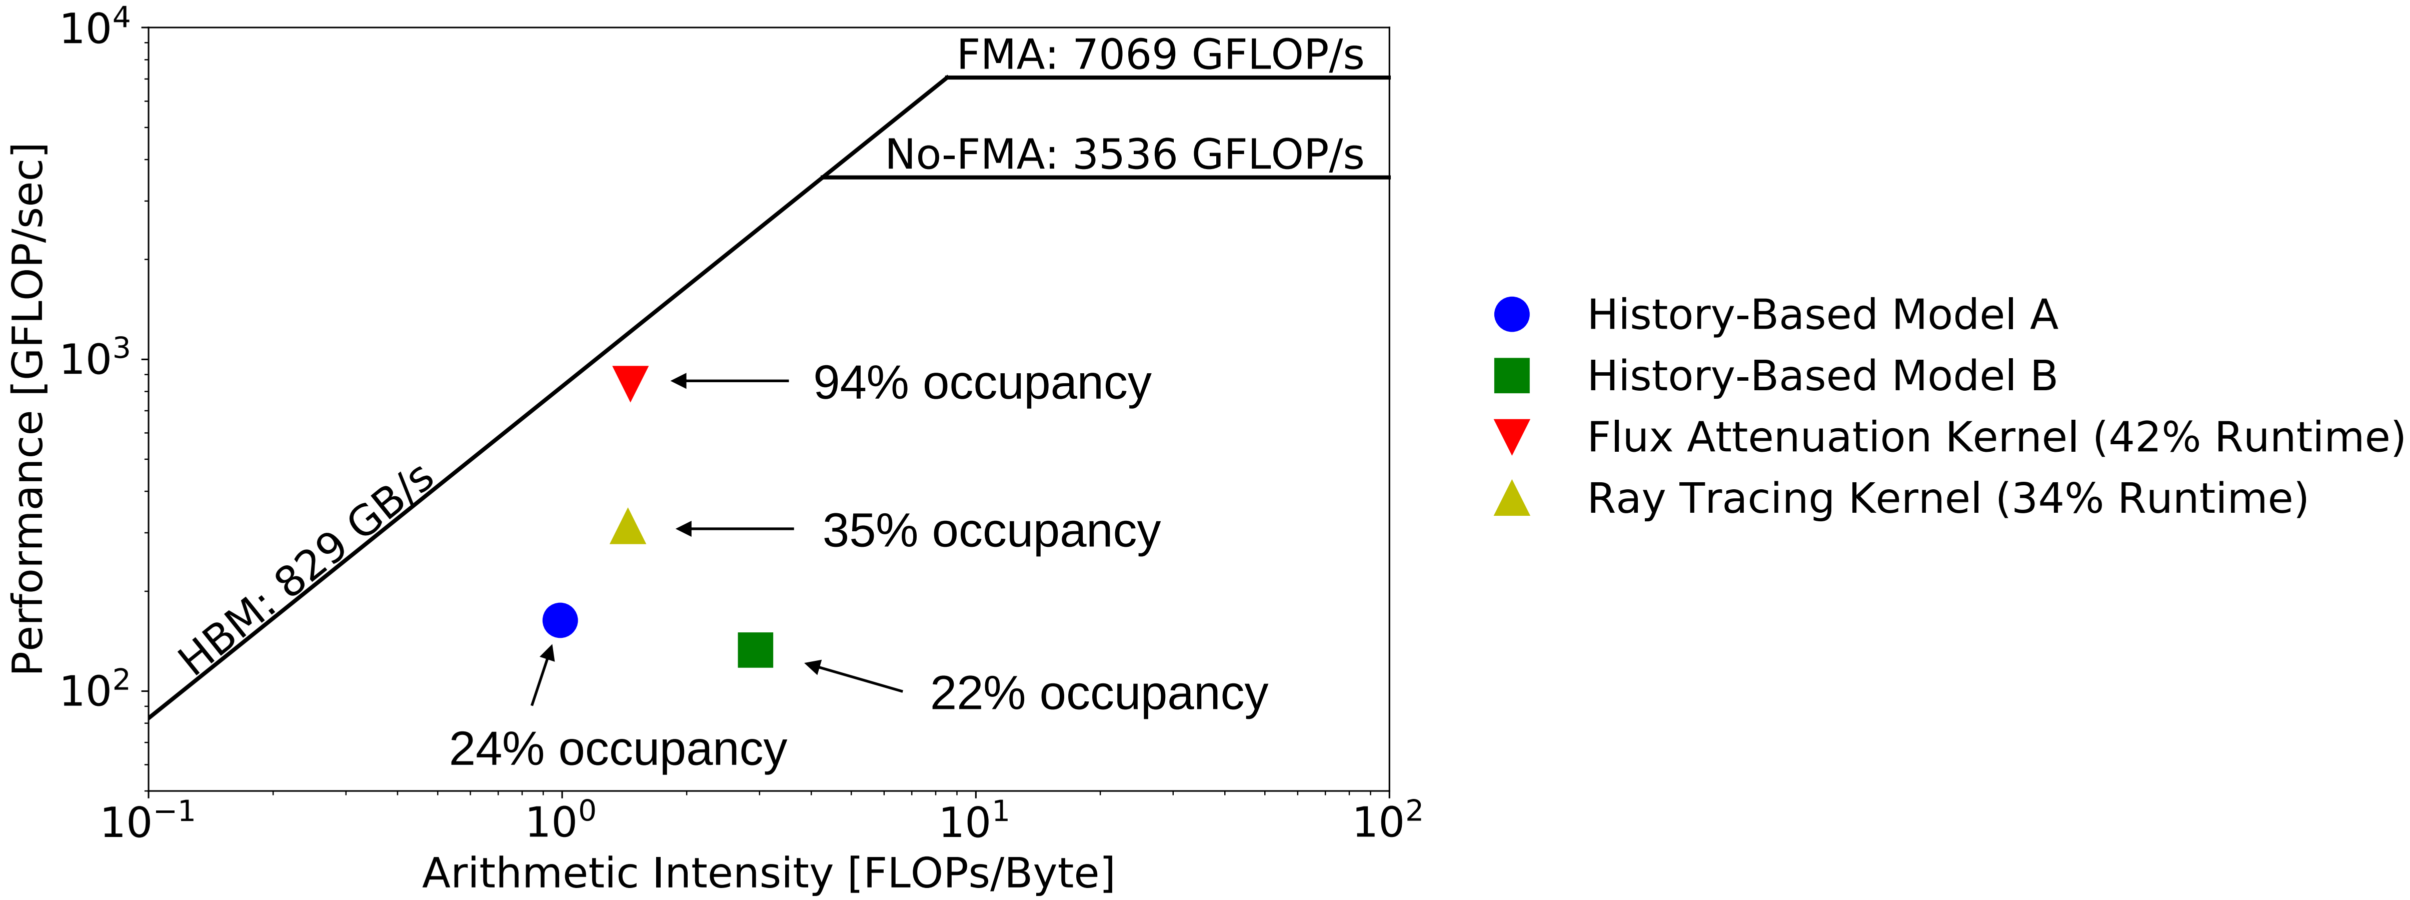

4) Choice of programming model

GPU-specific kernels

- Isolate the computationally-intensive parts of the code into CUDA/HIP/SYCL kernels.
- Refactoring the code to work well with the GPU is the majority of effort.

Loop pragma models

- Offload loops to GPU with OpenMP or OpenACC.
- Most common portability strategy for Fortran codes.

C++ abstractions

- Fully abstract loop execution and data management using advanced C++ features.
- Kokkos and RAJA developed by NNSA in response to increasing hardware diversity.

Co-design frameworks

- Design application with a specific motif to use common software components
- Depend on co-design code (e.g. CEED, AMReX) to implement key functions on GPU**.**

6) Managing software dependencies

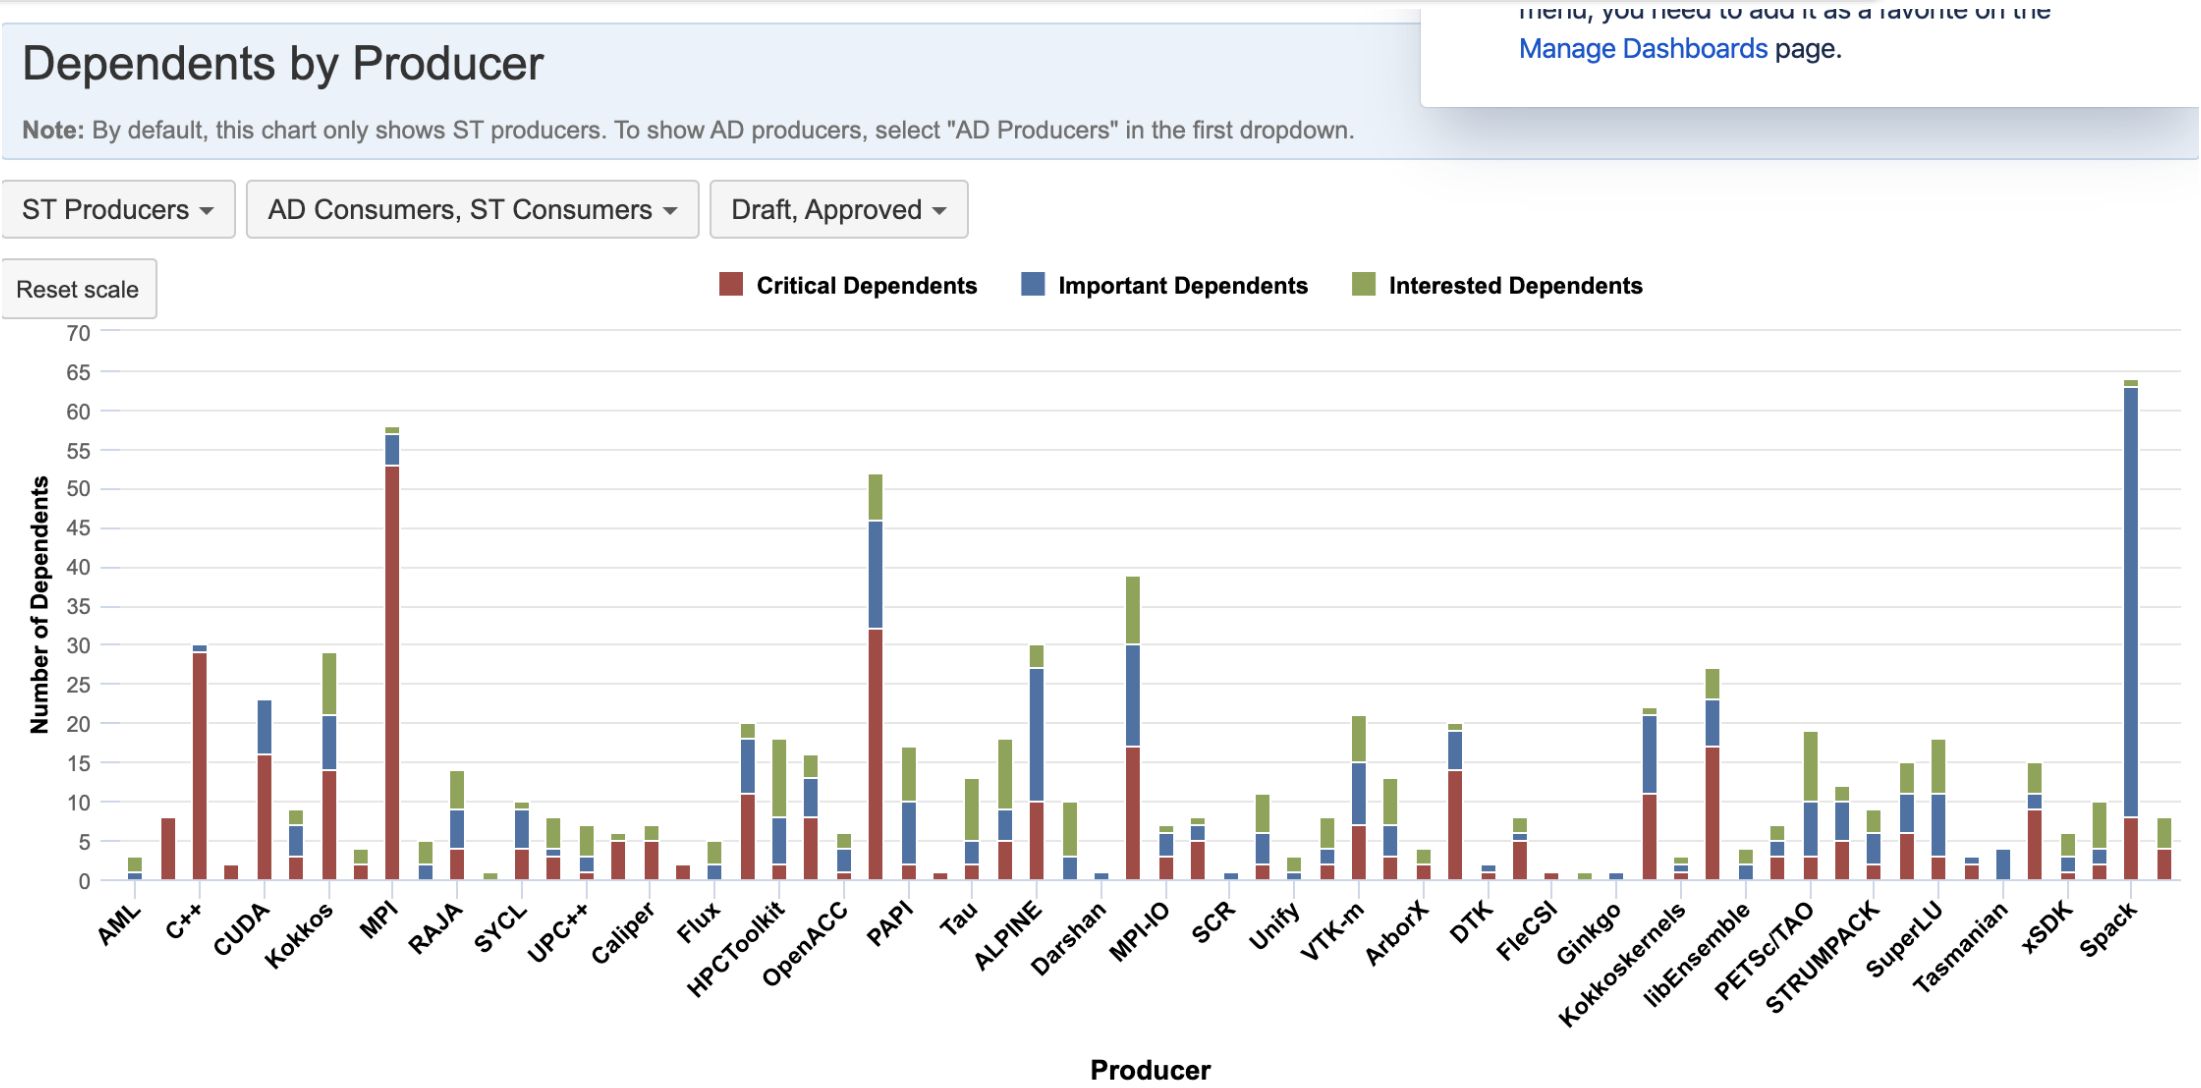

menu, you need to add it as a favorite on the
Manage Dashboards page.

Dependents by Producer

Note: By default, this chart only shows ST producers. To show AD producers, select "AD Producers" in the first dropdown.

6) Managing software dependencies

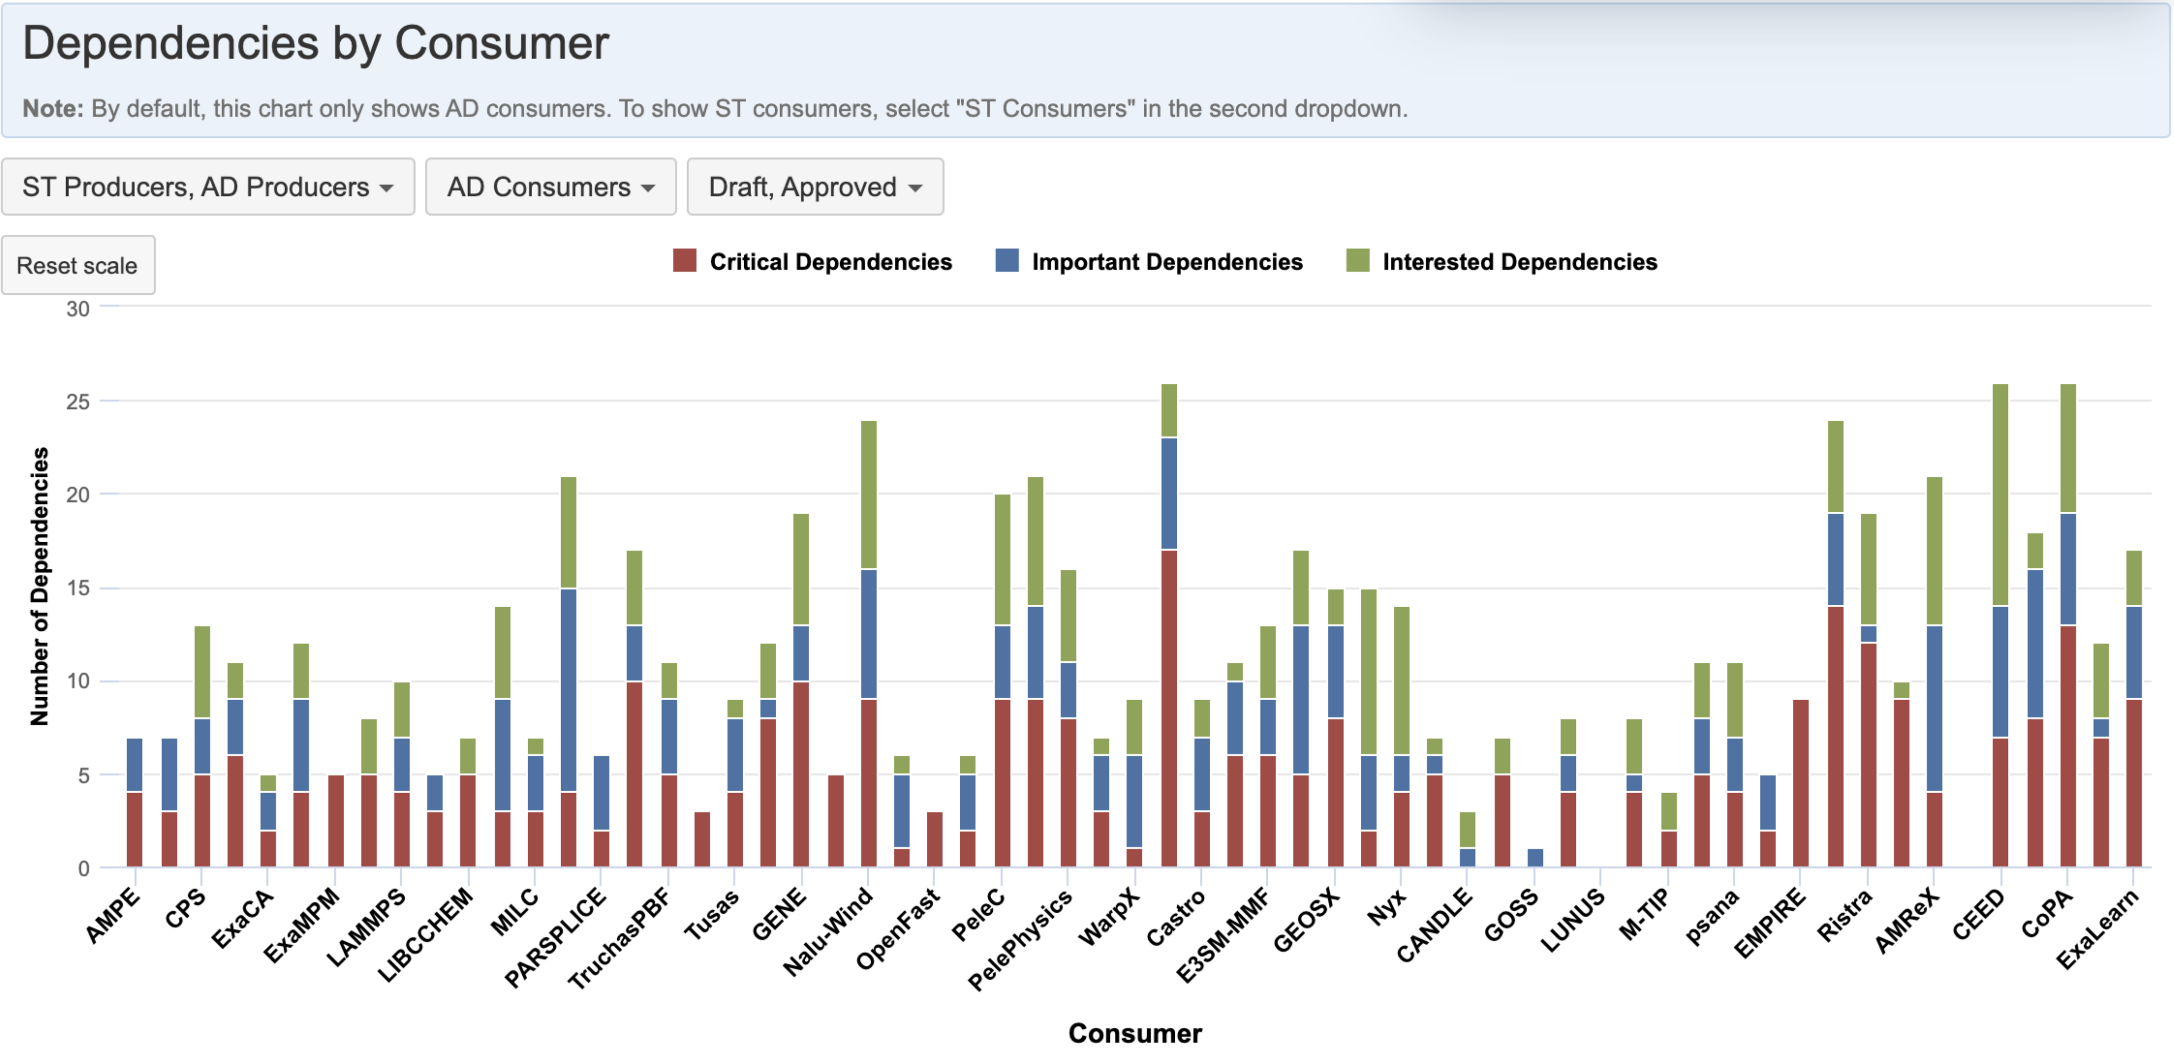

Dependencies by Consumer

Note: By default, this chart only shows AD consumers. To show ST consumers, select "ST Consumers" in the second dropdown.

AD codes use a mix of languages and programming models

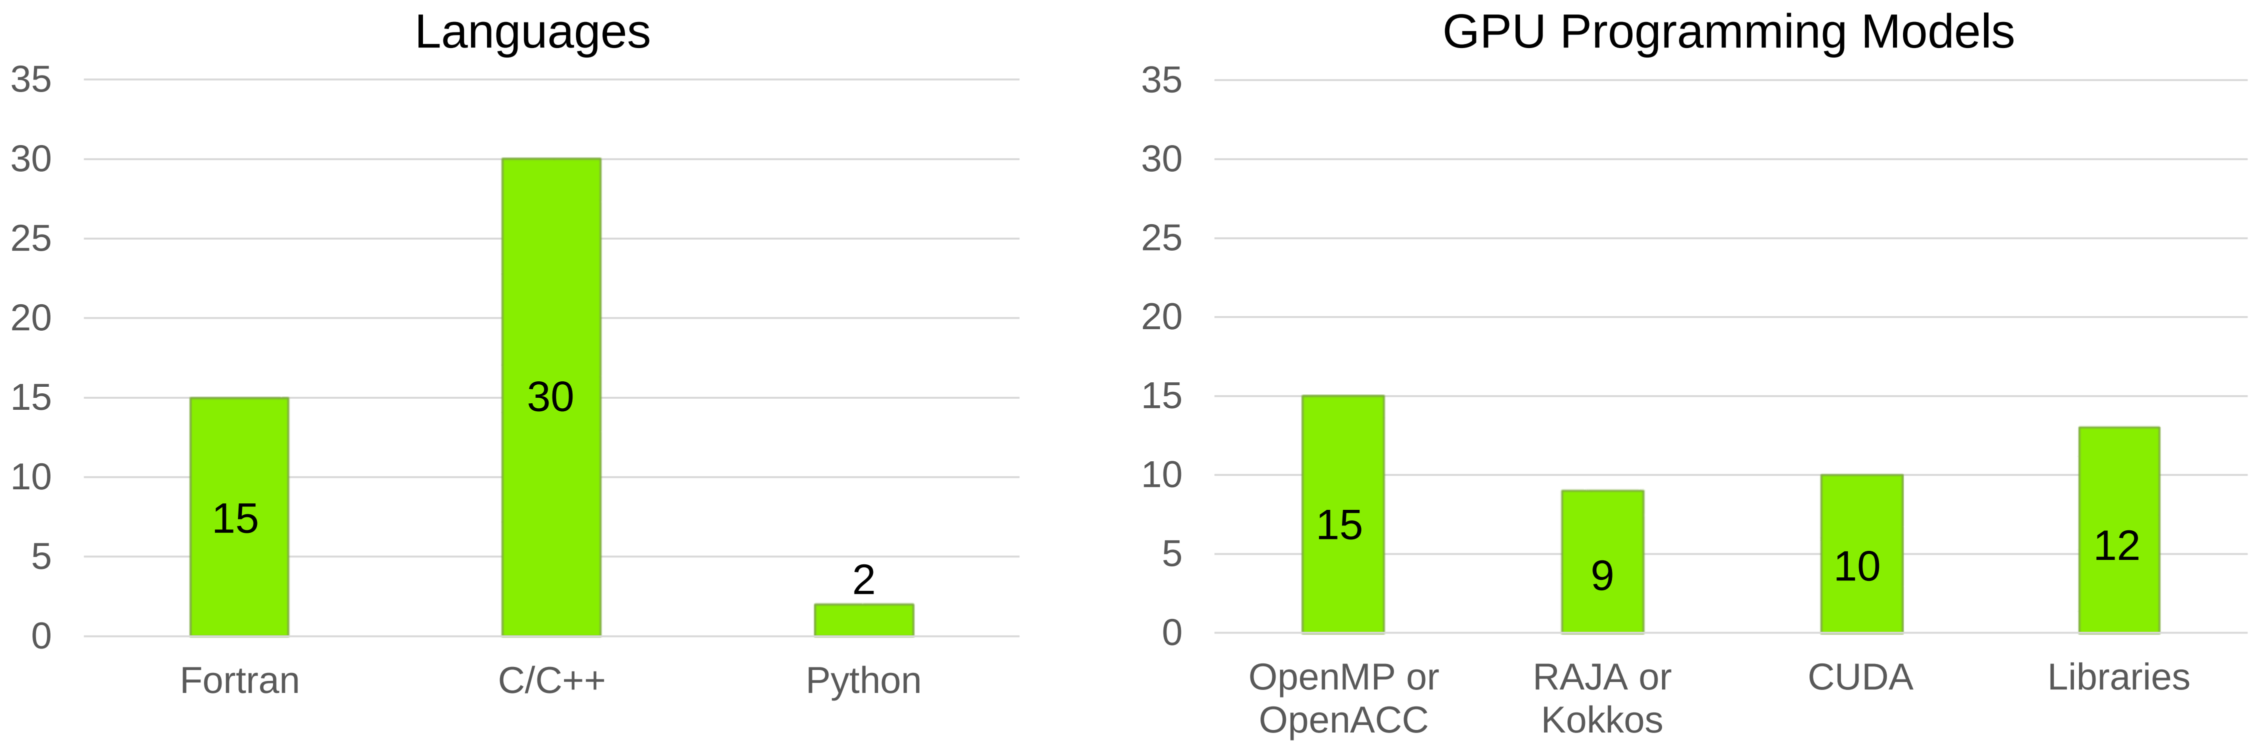

Many codes are still in flux, with quite a few still deciding on a final programming model. A few Fortran codes are being rewritten in C++, but most are not.

OpenMP/OpenACC: mostly Fortran users

Application Project	Code	Main Language	GPU Programming Model
ExaStar	FLASH	Fortran	OpenMP
ExaStar	CASTRO	Fortran, C++	OpenMP, OpenACC
E3SM-MMF	E3SM	Fortran	OpenACC, moving to OpenMP
Combustion-PELE	PeleC	Fortran	CUDA, OpenACC
Combustion-PELE	PeleLM	Fortran	CUDA, OpenACC
ExaSMR	Nek5000	Fortran	OpenACC
ExaSMR	OpenMC	C++	OpenMP, OpenCL or SYCL
WDMApp	GENE	Fortran	OpenMP
WDMApp	GEM	Fortran	OpenACC
WDMApp	XGC	Fortran	OpenMP, OpenACC
ExaBiome	GOTTCHA	C++	OpenMP, HIP, SYCL
ExaBiome	HipMCL	C++	OpenMP, HIP, SYCL
QMCPACK	QMCPACK	C++	OpenMP
ExaAM	MEUMAPPS-SS	Fortran	OpenMP, OpenACC
ExaAM	Diablo	Fortran	OpenMP

COVID-19 R&D in AD

- Change in scope requires ECP and DOE approval.
- Formal tracking of costs/scope
- Discourage sharp detour if can be avoided – syngeristic, fundamental R&D.

- ExaBiome
  - Performance evaluation, parallelization of the SpatialSim code
  - Exploring ancestral recombination and evolutionary origins of SARS-CoV-2 for vaccine development
- CANDLE
  - Workflow to identify small molecules that collectively target the entire SARS-CoV-2 proteome
  - identify protein targets, pockets, and drugs to combine; Identify proteins and binding pockets; accelerate search through billions of compounds
- ExaLearn
  - Apply surrogate and control techniques to emulate large-scale agent-based epidemiological models and explore dynamic (adaptive) intervention policies
  - Apply surrogate, design, inverse modeling capabilities to molecular drug design in partnership with CANDLE

Recent Highlights by Category

Performance Improvements

- **WarpX** (Jean-Luc Vay, LBL): new FOM measurement using 4,263 nodes (out of 4,608) of Summit. The new FOM is now **54 times the baseline** FOM (measured on 6,625 KNL nodes, out of 9,688), when extrapolating both FOM values to the full machines access to discrete AMD and Intel GPUs that are likely the foundation of their custom exascale accelerators

- **ExaSky** (Salman Habib, ANL): new “GPU-resident” variant of the HACC code’s first order CRK (Conservative Reproducing Kernel)-SPH hydrodynamic solver, designed to efficiently utilize accelerators, and maintain load balancing across millions of processors. Compared to the heavily optimized tree-based algorithms previously designed for CPU systems, the new solver achieves **8-12x performance improvements of the computationally demanding hydro solvers**

- **CANDLE** (Rick Stevens, ANL): new FOM calculation, showing **significant performance improvements after reducing memory usage of the P3B4 model**, which allowed for restructuring the model to improve data motion and expose additional parallelism during training. As a result of this restructuring, P3B4’s GPU utilization was improved on the NVIDIA V100s on Summit and recorded an FOM which is a significant improvement over the previously reported values.

Capability Demonstration

- **ExaBiome** (Kathy Yelick, LBL): developed **an experiment to demonstrate measurable advantages of co-assembly over multi-assembly**, including improved assembly of low depth (e.g., 5x depth) genomes (80% for co-assembly vs 5% for multi-assembly). Also demonstrated was the increased detection of genomes in real data (50% more genomes overall, with 4x more of high quality), improved contiguity, and reduced error rates.

- **EQSIM** (David McCallen, LBL) : carried out a **validation exercise for the coupling of the regional-scale geophysics finite difference wave propagation code SW4** with the structural / soil system finite element codes ESSI and Opensees. The coupling is accomplished through the Domain Reduction Method (DRM) and the intercode comparisons provided a validation of the implementation of the DRM. The validation exercise demonstrated that the ground motions created with an SW4 simulations exactly matched the ground motions generated with SW4 with an embedded soil island grid

Code Release

- **CEED** (Tzanio Kolev, LLNL): **released version 4.1 of the MFEM finite element library,** https://mfem.org. New features in the 4.1 release include: improved GPU capabilities including support for HIP, libCEED, Umpire, debugging and faster multi-GPU MPI communications; GPU acceleration in many additional examples, finite element and linear algebra kernels; many meshing, discretization and matrix-free algorithm improvements; ParaView, GSLIB, HiOp and Ginkgo support; 18 new examples + miniapps; significantly improved testing; and a new BSD license. More details can be found at
- **Proxy Applications** (Dave Richards, LLNLL): **released version 3.0 of the ECP Proxy App Suite at** proxyapps.exascaleproject.or**g**. The new release replaces the CANDLE benchmarks with a new miniGAN proxy app and updates the versions of existing proxies. This release also highlights selected proxies that are likely to be of significant interest to the ECP community. For example, AMD has released HIP versions of SW4lite, Quicksilver, and PENNANT. miniGAN is our first attempt to explore new proxies for Machine Learning (ML)
- **CODAR** (Ian Foster, ANL)**: released version 0.2.0 of the Feature Tracking Kit (FTK)** that incorporates start-of-the-art topological, statistical, and deep learning feature tracking algorithms for scientific applications. The FTK is scalable and thus enables in situ feature tracking with the simulations that run on today’s and future leadership computing facilities. This release includes new optimizations for tracking critical points in parallel with MPI and CUDA-based acceleration of these algorithms. A ParaView plugin that tracks minima, maxima, and saddle points in two dimensional scalar fields is developed and included in the release.

New Model or Algorithm

- **ExaFeL** (Amedeo Perazzo, SLAC): published a paper applying **pixel-level X-ray tracing to the data reduction step of protein crystallography diffraction experiments** for X-ray Free Electron Laser light sources. This is a highly anticipated development, because it promises increased accuracy in the measurement of small atomic structural details that are critical for understanding chemistry. The paper shows (in simulation) that the new method is sensitive enough to locate even a single electron at a metal atom within a protein.

- **Pele** (Jackie Chen, SNL): **Completed an initial GPU-portable multi-regime spray impingement**. Modeling of spray impingement upon a piston or cylinder surface is important for hydrocarbon emission and soot predictions in simulations of internal combustion engines with direct injections.

Next Steps

- Continue pushing performance envelope and testing on Summit, Sierra, and similar
- Work closely with ST to manage timeline and requirements for software dependencies
- Explore deeply and downselect of exascale programming model(s), including push/pull with vendors on compilers
- Develop and test new gpu-resident physics models for KPP-2 applications
- Understand key performance issues for initial target exascale architecture