The Exascale Computing Project

Carleton DeTar (University of Utah)

USQCD All-Hands Meeting
April 27, 2019
Brookhaven National Laboratory

EXASCALE COMPUTING PROJECT

Exascale Timeline

- Pre-exascale
  - ORNL Summit (pre-Exascale) already in production
  - TACC Frontera (pre-Exascale) in production mid 2019
  - NERSC Perlmutter (pre-Exascale) early 2020
- Exascale
  - ALCF Aurora (Cray + Intel Xe)
    - Early Science: late 2021 or early 2022
    - Production mid 2022
  - ORNL Frontier expected 2022

OLCF Frontier

- Exascale
  - ORNL Frontier expected 2022
  - 50 - 100 X Titan

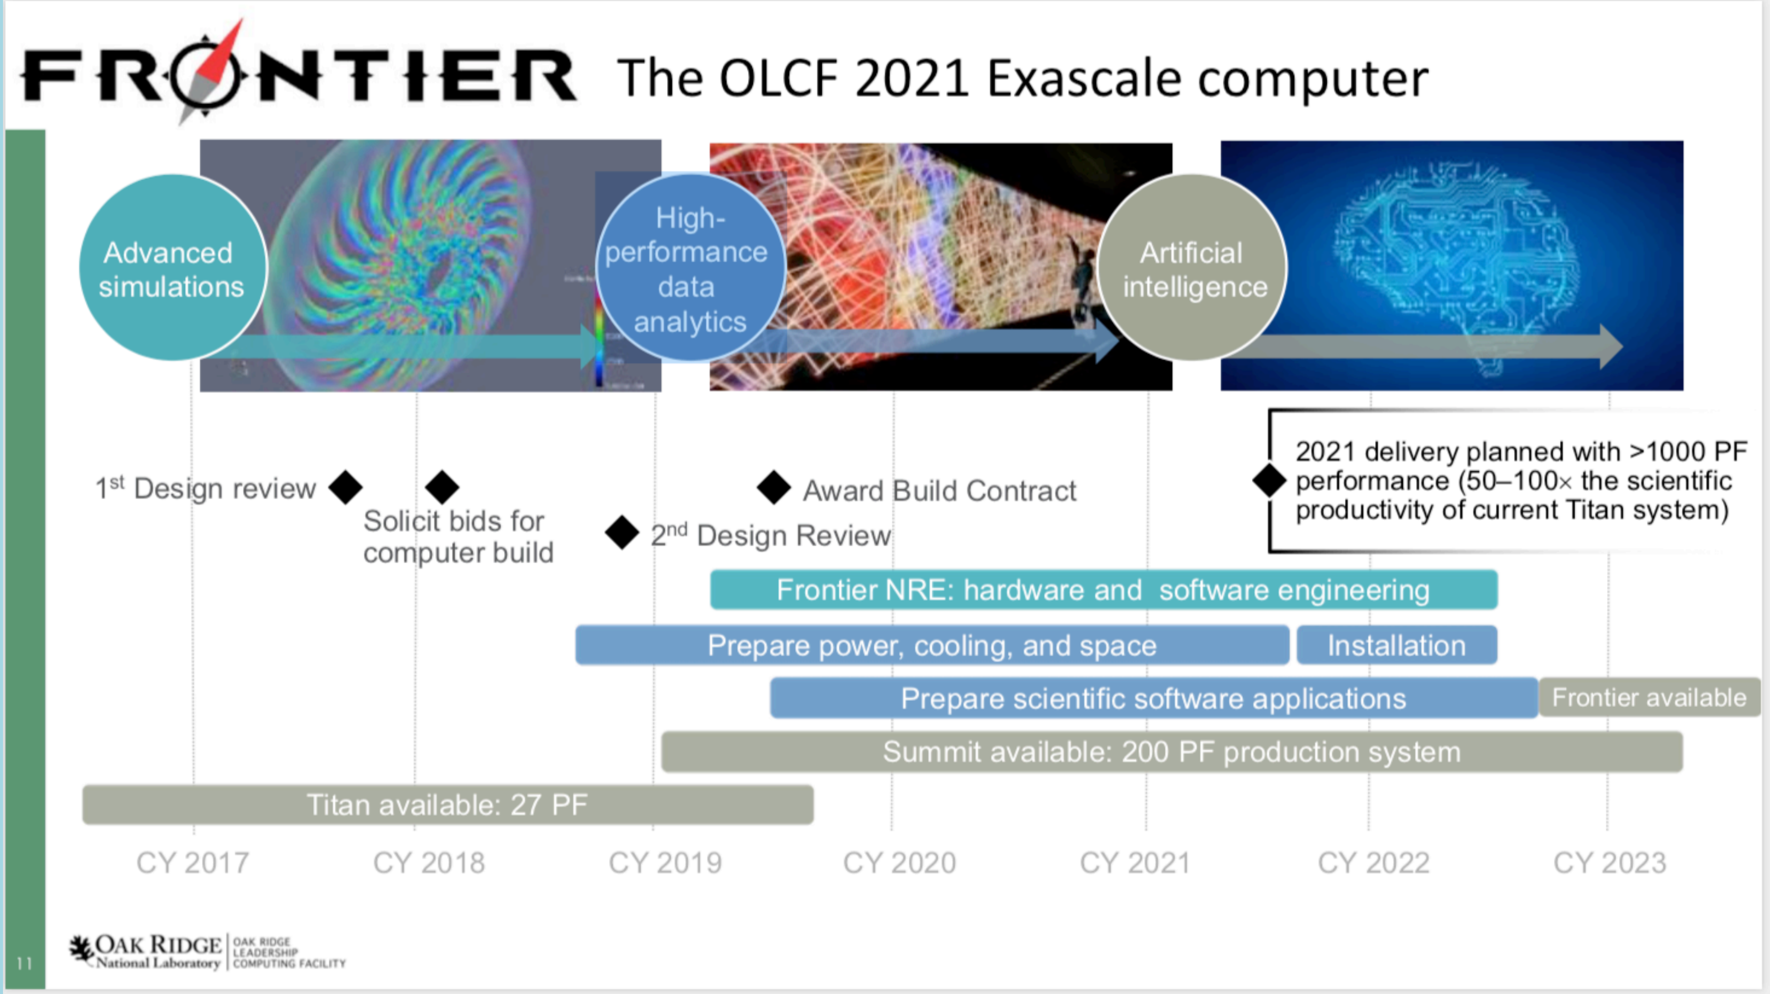

Exascale Computing Project Organization

- Joint DOE and NNSA project
- ORNL lead institution
  - Various DOE labs manage parts of the project
- Three sets of teams
  - AD Application Development (25)
    - 6 Chemistry and Materials applications
      - Lattice QCD
    - 6 Energy Applications
    - 5 Earth and Space Sciences applications
    - 5 Data Analytics and Optimization applications
    - 3 National Security applications
  - ST Software Technology
    - Such as MPI, OpenMP, math libraries, etc
  - HI Hardware Integration
    - Hardware design

ECP EXASCALE COMPUTING PROJECT

ECP Participating Institutions

A USQCD project
with significant outside participation

USQCD US Lattice Quantum Chromodynamics

- Fermilab
- BNL
- JLab
- ALCF
- Boston U
- Columbia U
- SUNY Stony Brook
- William and Mary U
- Indiana U
- U Illinois UC
- U Utah
- **Edinburgh U**
- **NVIDIA**

Approximate total FTE (15.5 funded + 5.1 unfunded)

USQCD Software Readiness

- New architectures present new computational challenges
  - Increased use of multicore (e.g. Intel Sky Lake, Crystal Lake, AMD EPYC) and GPU-like accelerators require new programming paradigms for maximum efficiency
  - Nodes are becoming fatter requiring more topology awareness in communications
  - Communication bottlenecks are becoming more severe
- Our physics problems are more demanding, requiring new attention to algorithms
- Our LatticeQCD Exascale project is intended to advance Lattice QCD software technology for the benefit of all of USQCD

Exascale Challenges for QCD

Critical slowing down

With decreasing lattice spacing the cost of generating a decorrelated gauge configuration grows due to topology freezing.

Communication bottlenecks

Internode communication may substantially reduce performance on exascale architectures.

Complexity of nuclear matrix elements

The Wick contraction complexity grows very rapidly with increasing atomic number.

Lattice QCD Project Organization

- Project PI: Andreas Kronfeld
- Software task (Carleton DeTar)
  - Develop exascale optimization strategies for our code bases, including adopting and incorporating a new data-parallel framework for lattice QCD calculations.
  - Participants: Utah, Argonne, Columbia, BNL, JLab, Edinburgh, NVIDIA
- Solver task: sparse matrix problems (Rich Brower)
  - Identify/develop communications-avoiding algorithms, especialy for solvers
  - Develop multi-scale algorithms, particularly adaptive multigrid.
  - Participants: Boston U, William & Mary, Illinois, NVIDIA, BNL, Fermilab
- Critical-slowing-down task (Norman Christ)
  - Identify/develop strategies for overcoming CSD in Markov sampling with decreasing lattice spacing.
  - Participants: Columbia U, BNL, Edinburgh, Argonne, William&Mary
- Matrix element (contractions) task (Robert Edwards)
  - Develop efficient strategies for dealing with the enormous combinatorial complexity of calculations involving light nuclei with multi-quark matrix elements
  - Participants: JLab, William & Mary
- Portability (Barbara Chapman, Meifeng Lin) (SOLLVE)
  - OpenACC, OpenMP standards, performance tools, e.g. Grid
  - Participants: Stony Brook

ECP EXASCALE COMPUTING PROJECT

ECP Challenge Problems and Performance Metrics

- ECP wants to see a > 50X improvement in the FOMs over Mira and Titan by a specified large fraction of the ECP applications.
- Physics we want to do *vs* how we quantify progress
- Physics we want to do
  - Support searches for new physics
  - Deeper understanding of hadron structure
  - Properties of light nuclei from first principles
  - Etc
- How we quantify progress
  - Three code bases/actions: CPS/DWF, MILC/HISQ, Chroma/Clover
  - Two benchmarks each: (1) lattice generation and (2) analysis
  - Progressively smaller lattice spacing
  - Our figure of merit (FOM) is the geometric mean of the rates.
- Our first Summit number: 7.5X

For more details

Software Meeting continues this afternoon