U.S. DEPARTMENT OF ENERGY | Office of Science

# Flash High-Energy Density Physics Co-Design Center


Director : Don Lamb

Associate Directors : Rajeev Thakur and Anshu Dubey

Argonne National Laboratory

University of Chicago


Anshu Dubey

IESP Workshop, San Francisco

April 2011







# Flash HEDP Co-Design Center Goals

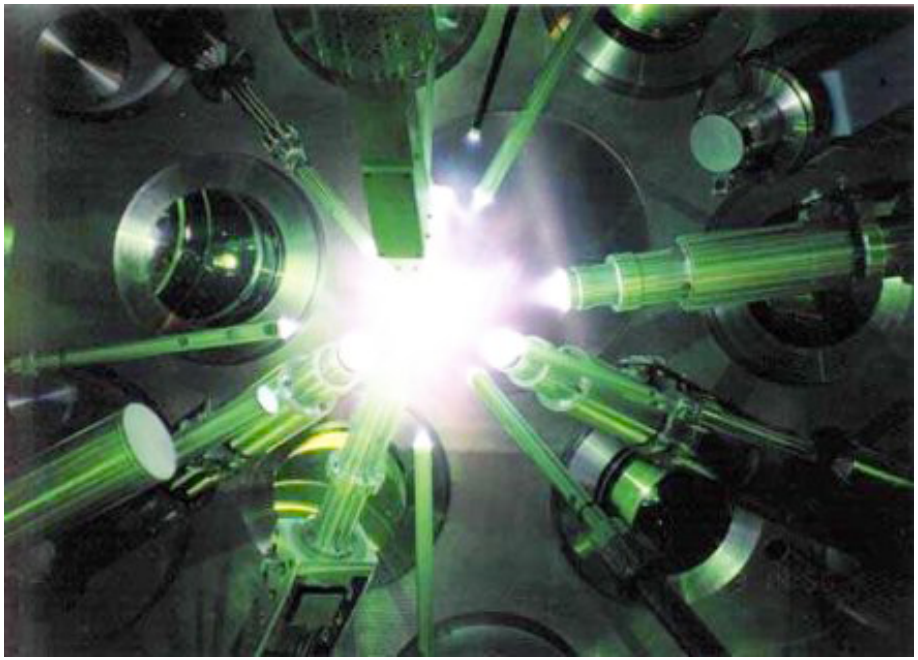

- Enable modeling and simulation of HEDP experiments on exascale computers
  - Allow scientists to address physical regimes and phenomena currently infeasible in the laboratory

- **Through a rigorously defined co-design structure and process, the Center will apply domain research to identify and tune interdependencies among the hardware, applied mathematics, and computer science, and the application**
- Incorporate insights from co-design into FLASH to ensure that it is a highly capable exascale code for the academic HEDP community

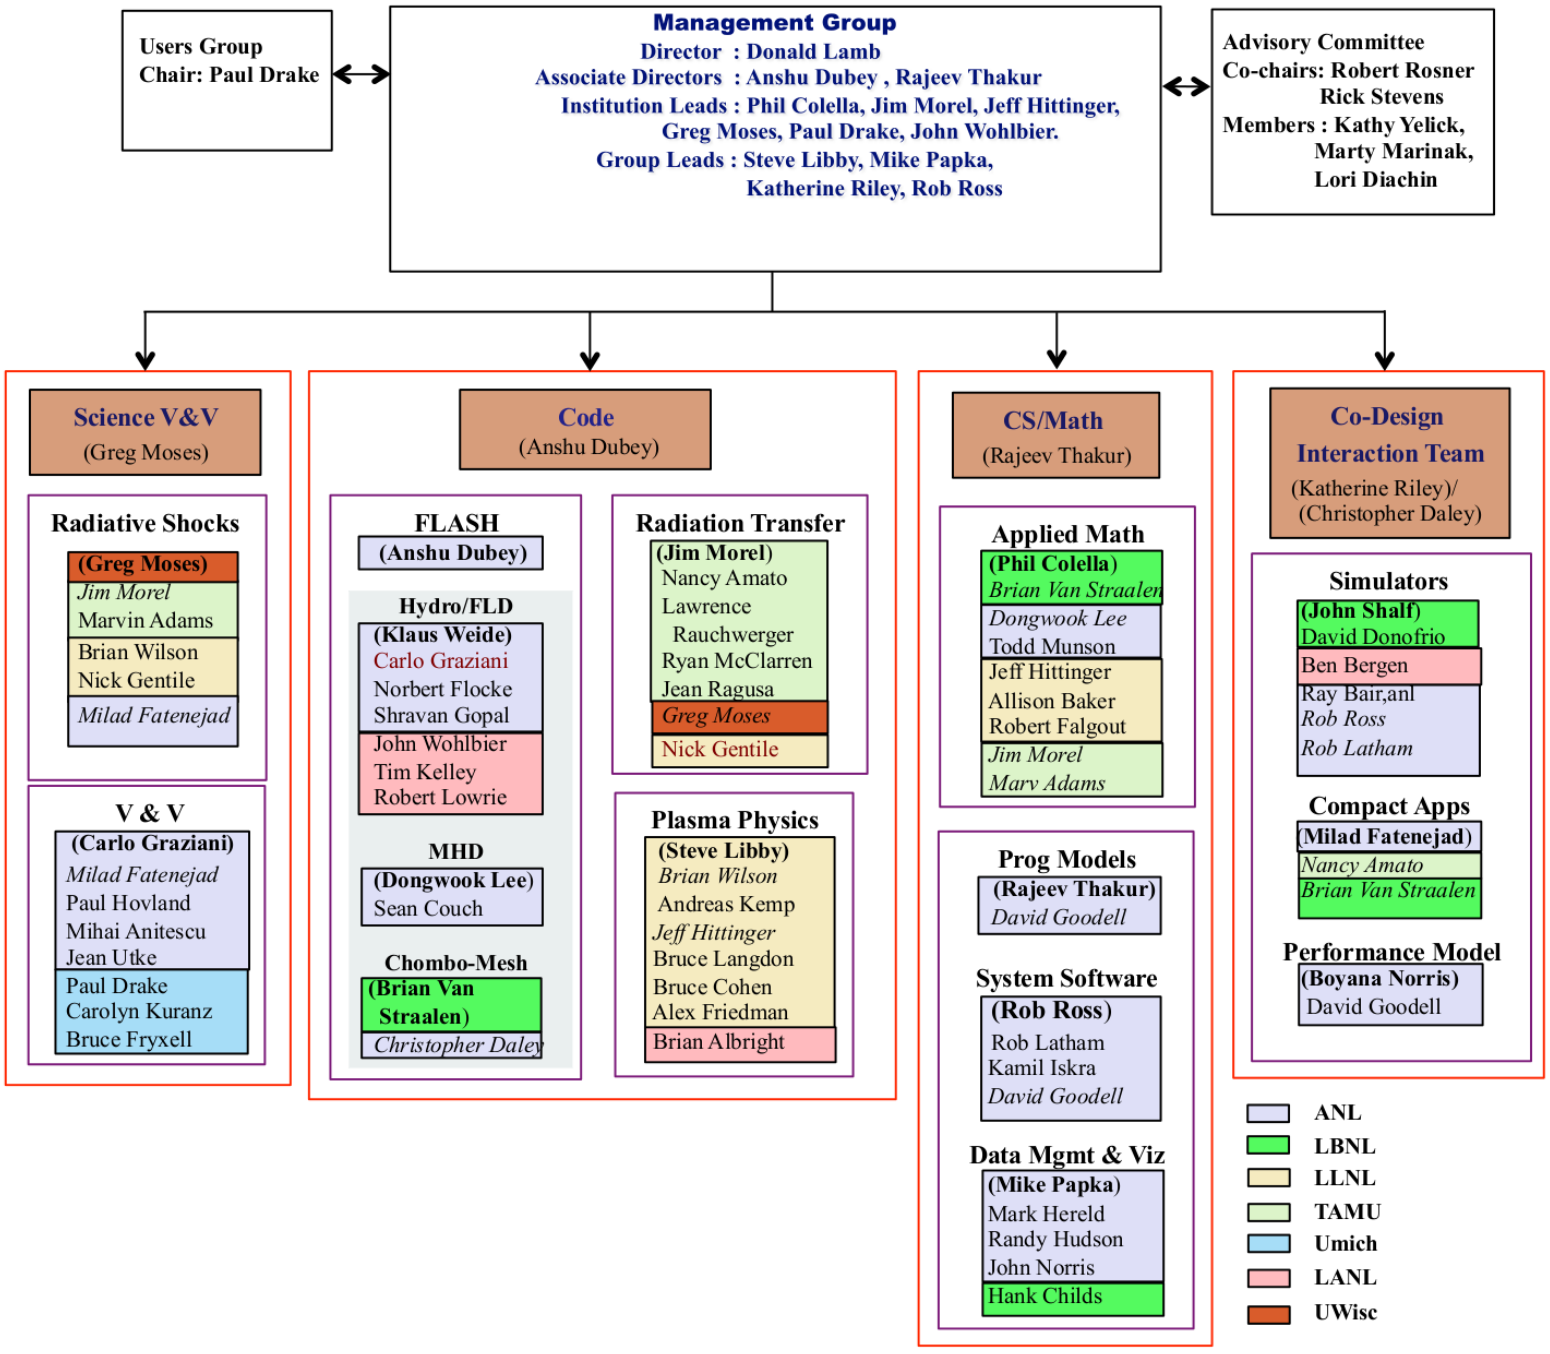
Users Group
Chair: Paul Drake
Management Group
Director : Donald Lamb
Associate Directors : Anshu Dubey , Rajeev Thakur
Institution Leads : Phil Colella, Jim Morel, Jeff Hittinger, Greg Moses, Paul Drake, John Wohlbier.
Group Leads : Steve Libby, Mike Papka, Katherine Riley, Rob Ross
Advisory Committee
Co-chairs: Robert Rosner, Rick Stevens
Members : Kathy Yelick, Marty Marinak, Lori Diachin
Science V&V
(Greg Moses)
Radiative Shocks
(Greg Moses)
Jim Morel
Marvin Adams
Brian Wilson
Nick Gentile
Milad Fatenejad
V & V
(Carlo Graziani)
Milad Fatenejad
Paul Hovland
Mihai Anitescu
Jean Utke
Paul Drake
Carolyn Kuranz
Bruce Fryxell
Code
(Anshu Dubey)
FLASH
(Anshu Dubey)
Hydro/FLD
(Klaus Weide)
Carlo Graziani
Norbert Flocke
Shravan Gopal
John Wohlbier
Tim Kelley
Robert Lowrie
MHD
(Dongwook Lee)
Sean Couch
Chombo-Mesh
(Brian Van Straalen)
Christopher Daley
Radiation Transfer
(Jim Morel)
Nancy Amato
Lawrence Rauchwerger
Ryan McClarren
Jean Ragusa
Greg Moses
Nick Gentile
Plasma Physics
(Steve Libby)
Brian Wilson
Andreas Kemp
Jeff Hittinger
Bruce Langdon
Bruce Cohen
Alex Friedman
Brian Albright
CS/Math
(Rajeev Thakur)
Applied Math
(Phil Colella)
Brian Van Straalen
Dongwook Lee
Todd Munson
Jeff Hittinger
Allison Baker
Robert Falgout
Jim Morel
Marv Adams
Prog Models
(Rajeev Thakur)
David Goodell
System Software
(Rob Ross)
Rob Latham
Kamil Iskra
David Goodell
Data Mgmt & Viz
(Mike Papka)
Mark Hereld
Randy Hudson
John Norris
Hank Childs
Co-Design
Interaction Team
(Katherine Riley)/
(Christopher Daley)
Simulators
(John Shalf)
David Donofrio
Ben Bergen
Ray Bair,anl
Rob Ross
Rob Latham
Compact Apps
(Milad Fatenejad)
Nancy Amato
Brian Van Straalen
Performance Model
(Boyana Norris)
David Goodell
ANL
LBNL
LLNL
TAMU
Umich
LANL
UWisc

# HPC System Timeline

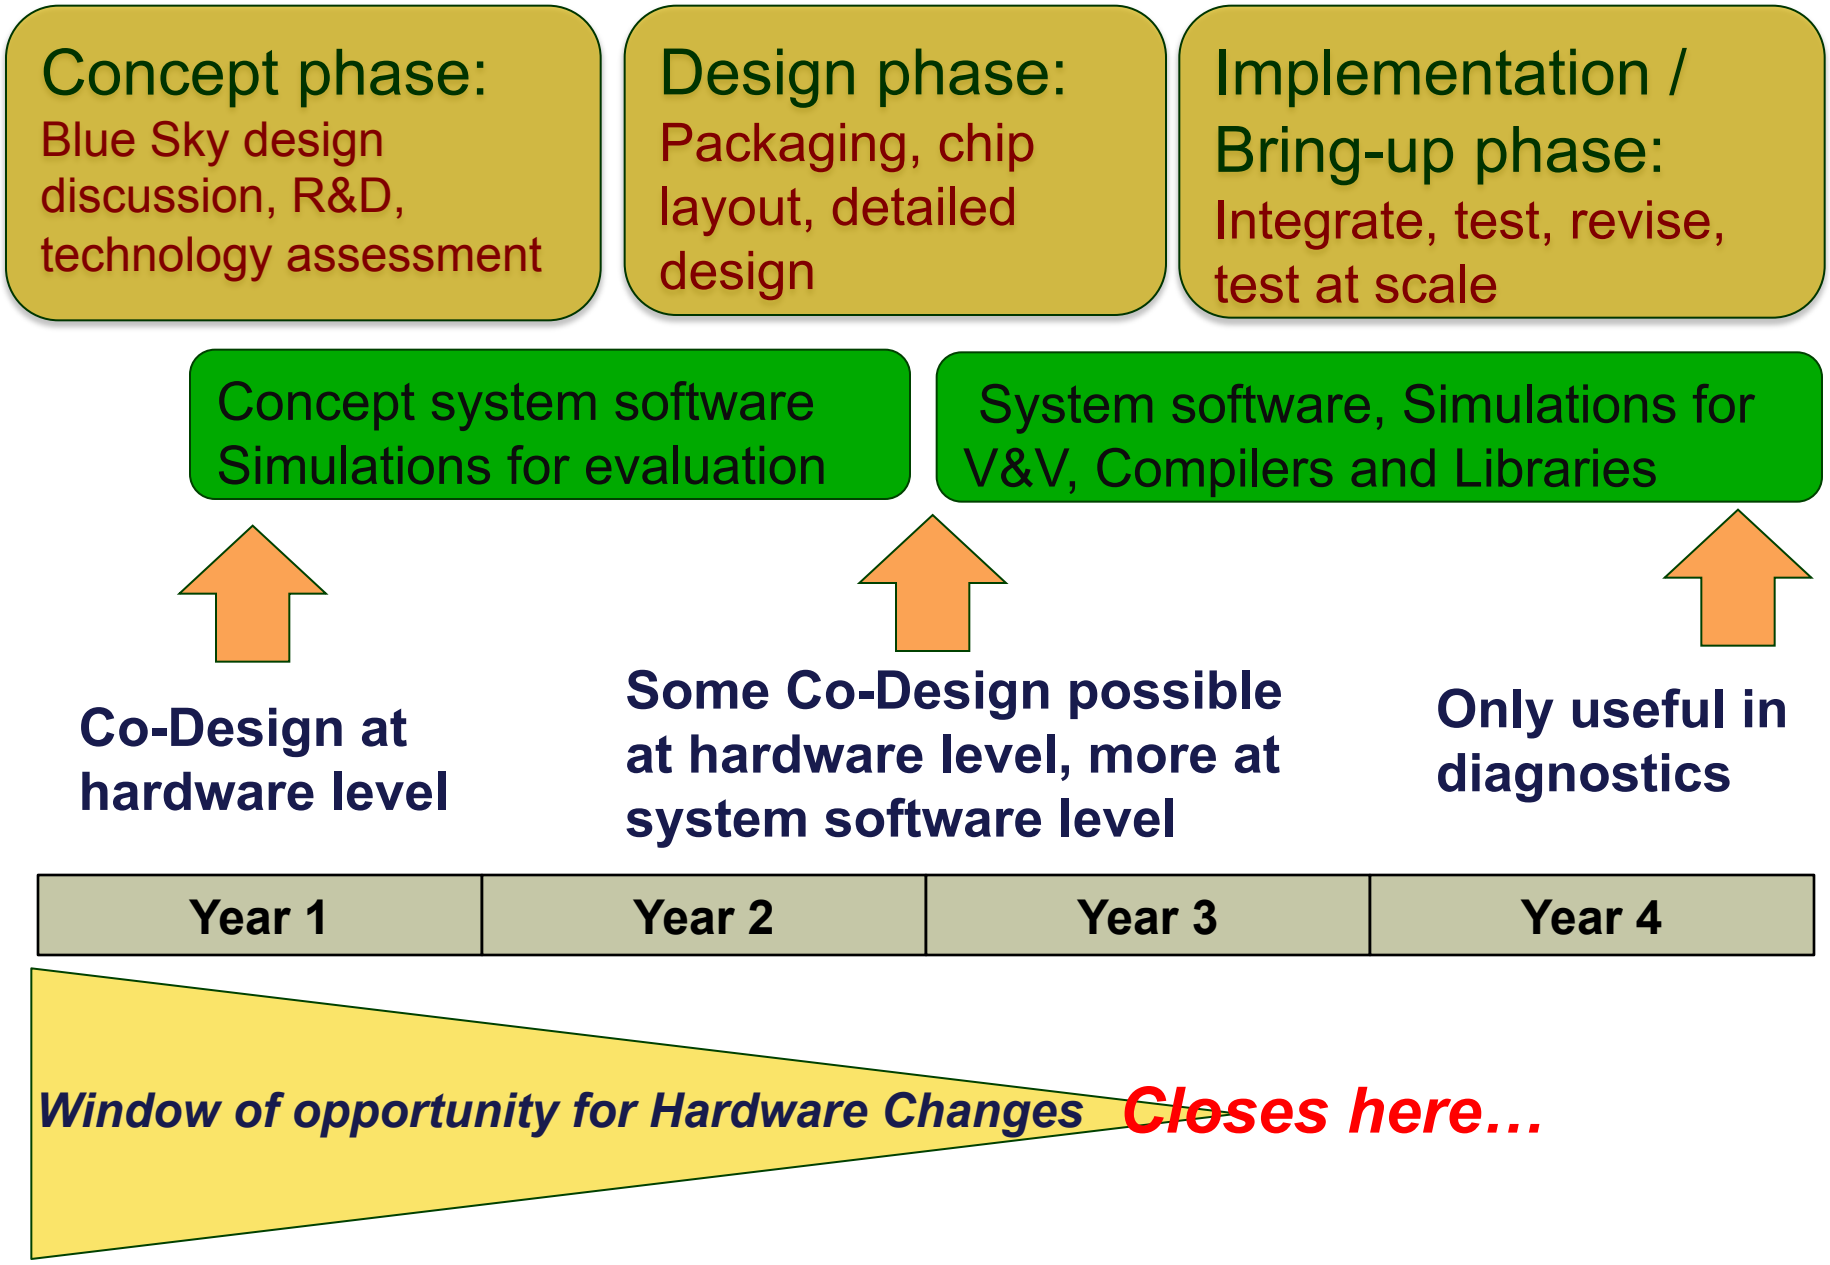

# Flash Co-design Interaction

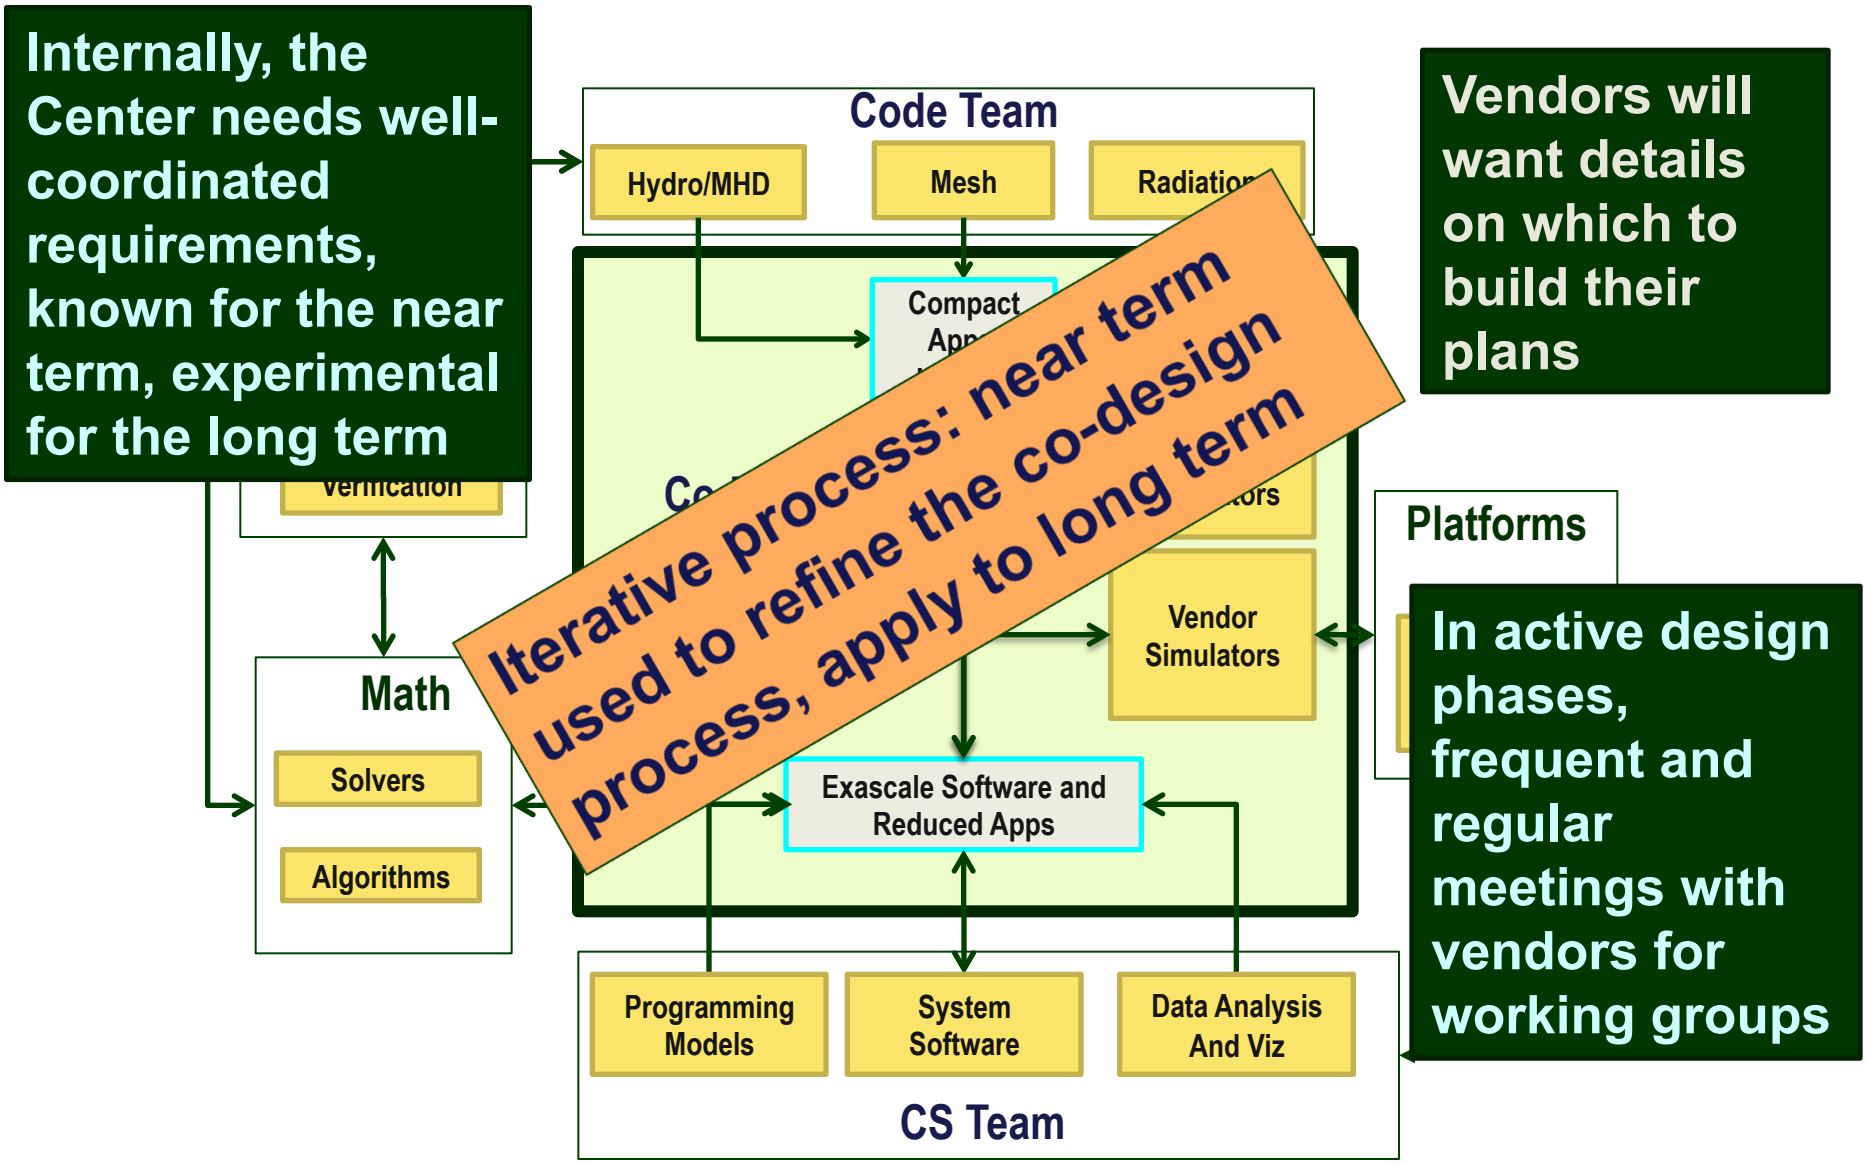

# Integration

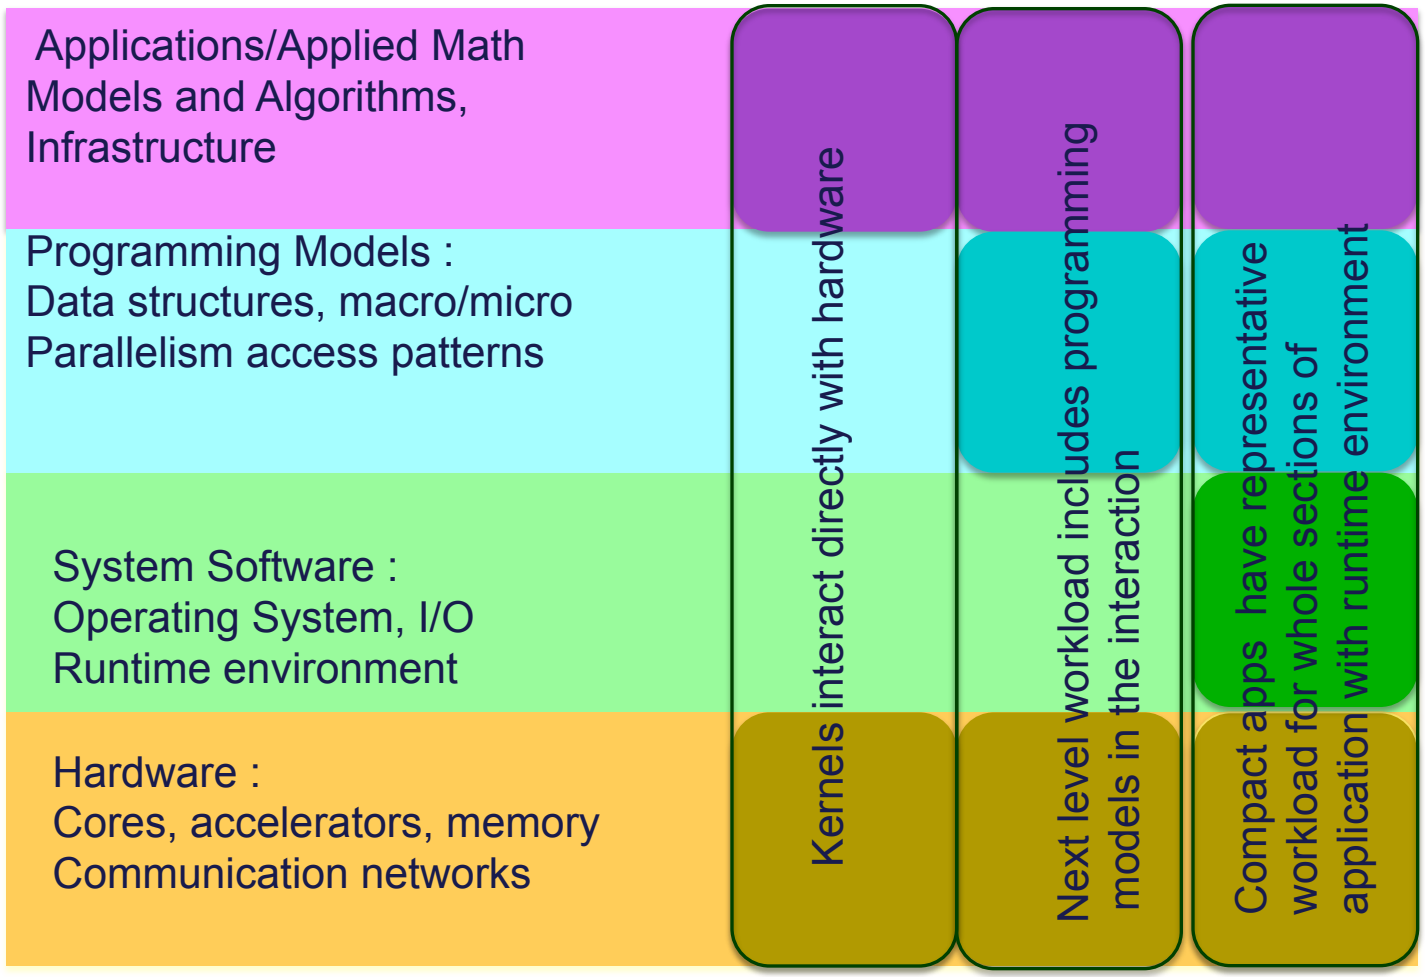

# Repository of Representative Workload

## Designated group for managing representative workload in the Co-Design Interaction Team

- Identify code components in the critical path
  - Infrastructure: meta-data and load balancing
  - Physics : memory intensive vs. computation intensive
- Identify code components that exercise co-design
  - IO, AMR, analysis

## Maintain a well documented repository of kernels, reduced and compact apps

- Available to vendor partners, exascale projects, and other collaborators
- Regularly updated based on feedback

# Examples of Vendor Interaction

- Flash an Early Science project for IBM BG/Q Mira
  - An official acceptance test
- Details under NDA, but one person working with IBM

- NVIDIA's HPC initiative has licensed FLASH for co-design
  - Multiphysics and multiscale
- Understand relatively easy changes they might make to CUDA GPU architecture that would allow FLASH to be more easily ported to it

- Because of regression testing, we understand representative workload
- Unit tests become reduced apps; comparison tests are compact/mini apps
- These can be easily adapted for vendor use (IBM already has the needed set; we are preparing one for NVIDIA)

Suitably stripped and documented code release licensed by vendors

# Inter-node Challenges

## Challenges

- Parallel IO
  - Analysis memory snapshot a large fraction of total system memory
- Higher degree of macro parallelism
  - Load balance
  - Meta-data handling
- Higher fidelity physics dictates greater coupling
  - Implicit/semi-implicit treatment

## Co-Design Opportunities

- Different approach through data staging
  - Critical vs. non-critical data
  - Combine with *in situ* analysis
- New parallel algorithms
  - Trade-off between duplication and communication
  - Possibly more hierarchy
- Investigate different class of numerical algorithms
  - Less deterministic

# Intra-Node and Resiliency Challenges

| Challenges | Co-Design Opportunities |
|---|---|
| **Intra-Node** | |
| ❑ Memory intensive computations<br>❑ Increasing limits on available memory per process<br>❑ Bigger working sets | ❑ Aggressive reuse of memory<br>❑ Distinguish between cores<br>❑ New algorithms<br>❑ Programming model |
| **Faults** | |
| ❑ Frequent failures<br>❑ Silent errors | ❑ Stochastic algorithms<br>❑ Redundancy |

# Code Maintenance and Co-Design

- Code verification and regression testing
  - Expect more non-determinism and async execution models to get performance and scalability
  - But to do regression testing without reproduci
  - Will study approaches to selectable d
    - Changes to compiler or run
    - Changes in algorit

Performance vs. Maintainability vs. Portability
Auto-tuning, code to code translation, annotations

- Tin science advances into the
  - t obsolescence of code modules)
  - et apps are living documents
  - Testing of new algorithms / implementation coming about because of new knowledge/insights

# Co-Design Needs from Application

- Greater encapsulation
  - Minimize common data
  - Maximize code sections that are re-entrant
  - Increase isolation between layers
  - Separate code functionalities such that different optimizations are applicable to different layers
- Minimize kernel dependency on programming models
- Expose optimization and fault tolerance possibilities
  - Be clearer about dependencies
  - Identify critical sections Vs the non critical sections
  - Define more compact working sets
- Explore more inherently robust alternative algorithms
  - Stochastic Vs deterministic

# Co-Design Needs from Hardware and Software

## During Co-Design

- Ability to express and implement ideas from the last slide
- Framework for testing ideas and experimentation
- Ability to realistically assess impact on algorithm and model choices
- Ability to evaluate direct influence of architecture decisions
- Metrics for performance evaluation
  - Deterministic performance engineering parameters
- Fault notification and recovery models
- Tuning parameters

# Co-Design Needs from Hardware and Software

## At Exascale

- Measurable and predictable performance
- Reliable results within quantified limits
- Retain code portability and performance
  - Standardized interfaces for common functionalities
  - Libraries and middleware
  - Auto-tuning or code to code translation
- Memory management
  - Memory bound application
- IO management
  - Large volumes of analysis data
  - Currently one snapshot roughly 1/10$^{th}$ of memory footprint
  - Analysis a judicious combination of in-situ and post processing