A look at trends in computational methods: the tie between machines and mission science

Aimee Hungerford

Los Alamos National Laboratory

40TH ANNIVERSARY LECTURE SERIES

LA-UR-21-23988

Discovery science happens in the intersection between machines & mission

- Broad trends in codes delivering "discovery science" for multiphysics applications
- Machines
  - Nic Lewis gave a nice overview of platform development
- Mission
  - Ray Juzaitis gave a nice overview of NNSA mission drivers throughout the ASCI and ASC eras
  - Examples from astrophysical transients

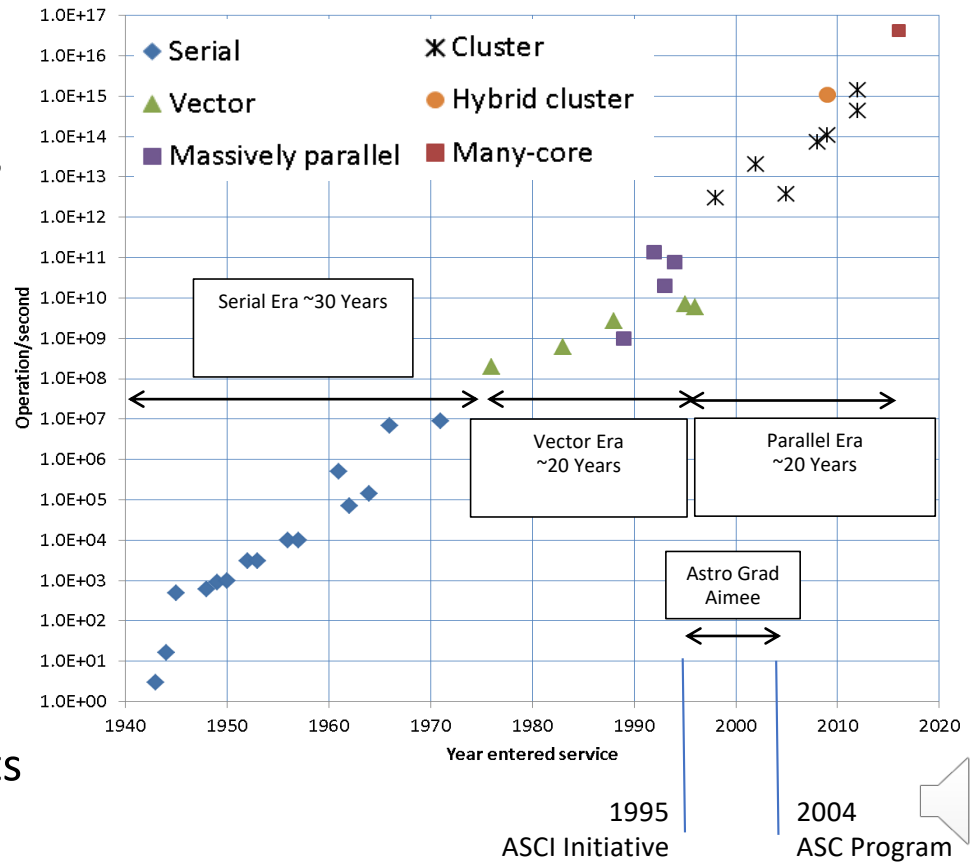

Pre-ASCI: Qualitative understanding from computation

1980's

Architectures

- Vector

Code technologies

- 1D supernova simulations aimed at understanding energetics, timescales
- What is key physics in driving supernova explosions?

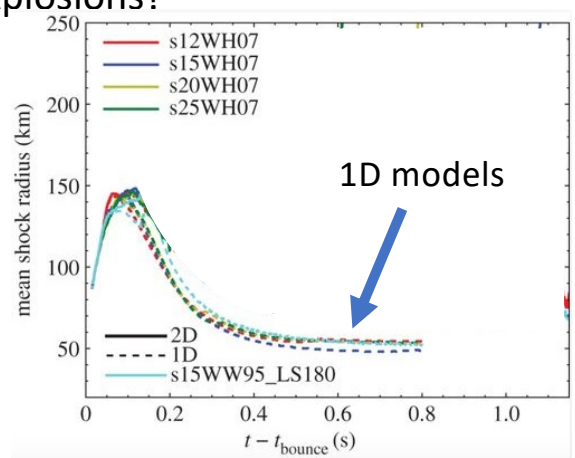

- Detailed radiation hydrodynamic, and microphysics in 1D does not show robust explosions

Mission/application drivers

- Blink comparison technology for finding supernova

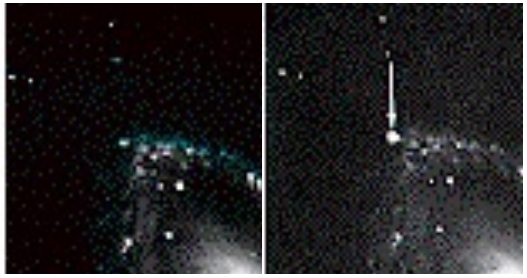

- Experimental program is the glue
- Computation for trends

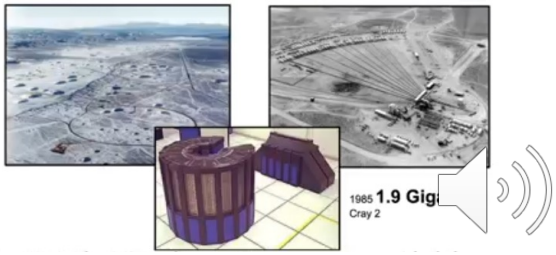

Pre-ASCI: Multi-dimension and multi-physics trends

1990's

Architectures

- MPP

Thinking Machine CM-5

- Supercomputer market crash
- Reduction of Gov't purchasing
- Rise of the PC

Code technologies

- More routine runs in multi-dimension
- Do we have the key physics yet?

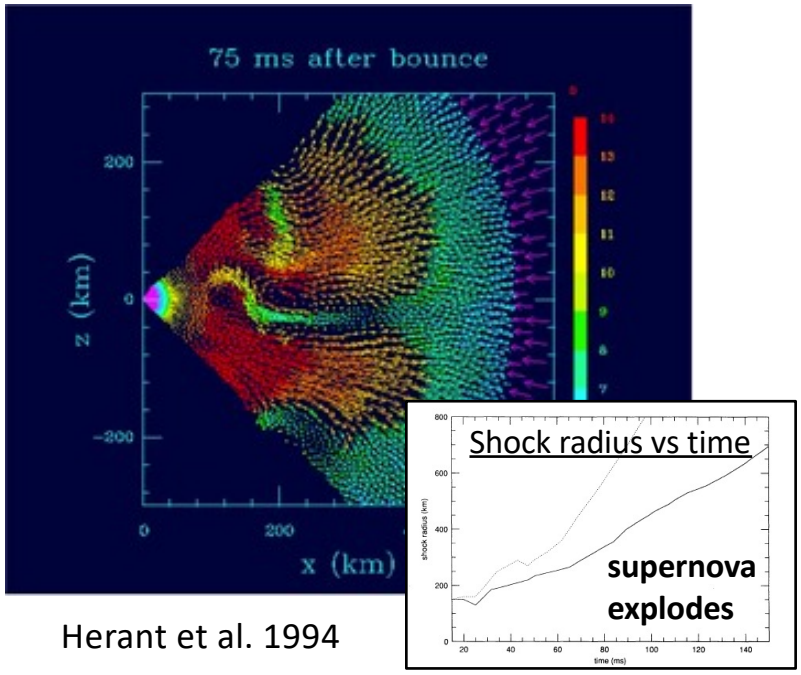

Herant et al. 1994

Mission/application drivers

- Supernova 1987A

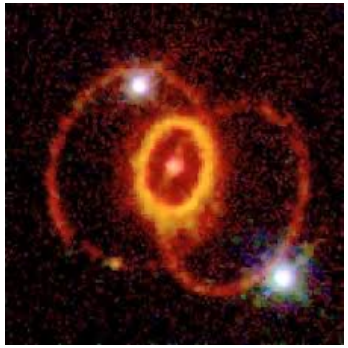

- CTBT → Stockpile Stewardship
- Mod/Sim becomes the glue

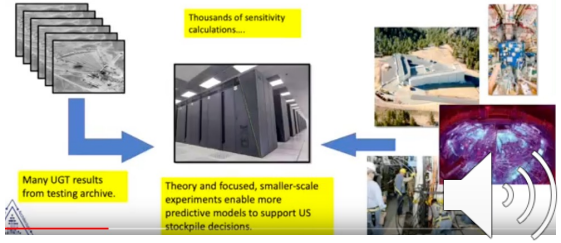

ASCI/ASC: Quantitative; predictive capability

~mid-90s-2015

Architectures

- Commodity cluster parallel
  - Economic driver
  - Long-lived stability

The Romans in their Decadence
Thomas Couture (1847)

from Randy Baker

- Clock speeds double every two years! (until early 2000's)
- Memories get larger!
- Vectors disappear!
- MPI prevails!
- Push for new codes

Code technologies

- MPI standard
- New MPI parallel multiphysics codes
  - Zeus-MP; ASCI Flash, CHIMERA
- How much do the details matter?
  - Add more!

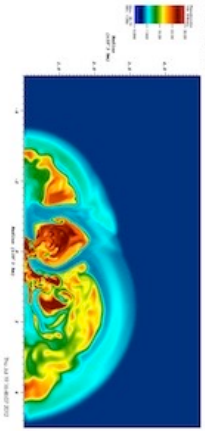

CHIMERA code ORNL

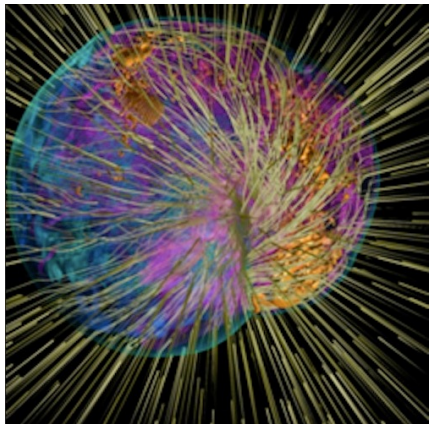

"Add more" mentality fit right in at the Labs!

Mission/application drivers

- Automated detections
- Statistics of a population
- Multi-messenger astronomy

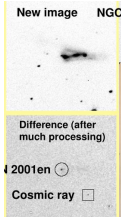

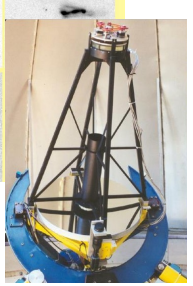

- Predictive Cap. Framework

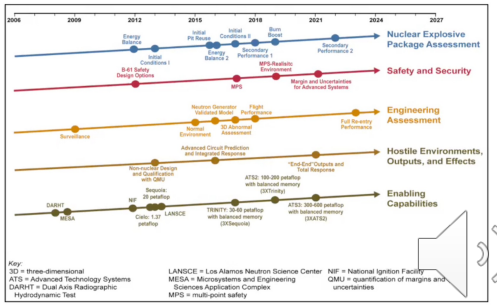

ASC: Abstractions, modularity, open source

~2010 - now Model form exploration, multi-scale models

Architectures

- Cluster parallel – diverse specialized hardware
  - GPU
  - *Component Tailoring*

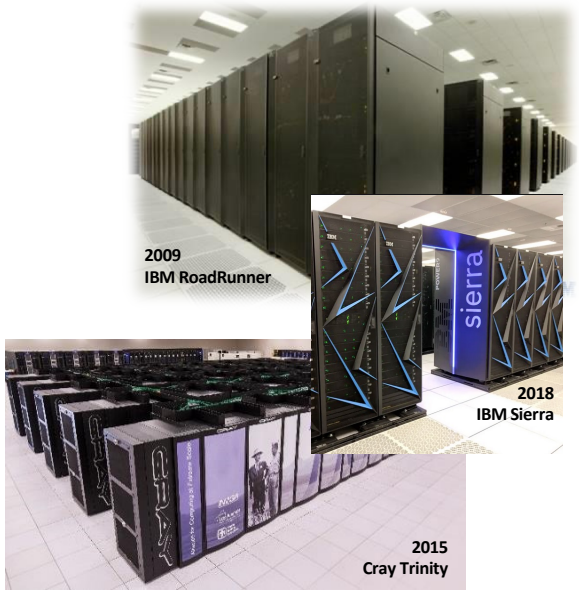

Code technologies

- Community codes; frameworks
  - Athena++, Gizmo, AMReX
- Performance portability abstractions
  - Node level: e.g. Kokkos …
  - System level: FleCSI …
- Beyond Euler & continuum approximation
  - Brute force resolution is not an affordable path to improving our understanding
  - Code agility to facilitate novel methods exploration is a priority

Mission/application drivers

- Social media…on demand data
  - Machine Learning
  - Cell phones & wearable tech

- Stockpile modernization
  - PCF → SCDS; pivot in focus with new scope

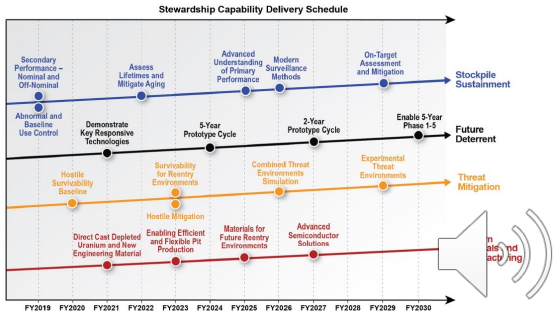

Highlight 2 LANL efforts aimed at addressing challenge of machine diversity & evolving mission

Next generation platforms (NGP)

We and Livermore have advanced a strategy that explores two paths enabled by commercial trends

Accelerated Architectures
Enabled by Industry-Driven Architectural Trends
Codesigned to Application Demands

Lawrence Livermore National Laboratory

Complex Simulation Architectures
Leveraging Component-Level
Customization Tailored for Application

Los Alamos NATIONAL LABORATORY

The challenge with this path is in finding an economical way to shape architectures.

NVIDIA partnership as a path to success

ATDM Next generation code (Ristra/NGC)

Design of a high level abstraction layer FleCSI to aid in efficient porting, and enable physics agility

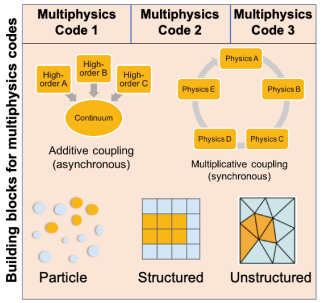

Abstraction layers
Facilitate diverse implementations?

Abstraction layers
Automate the port?

ATDM investment initial success in L1 milestone

Insulate from architectural diversity when possible
Codesign to optimize for our most complex workflows

NGC: Ristra design targets flexibility in a volatile future

Separation of concerns between physics expression and CS execution

Ever changing scientific questions & Rapidly changing HPC hardware

- Ristra is creating a ***flexible and agile*** toolkit for next-generation multi-physics
- *FleCSI* abstraction layer separates concerns of physics and CS developers
- *Portage* remap and link library connects meshes, packages, and codes

Current codes are

- Synchronous; Static; Difficult to extend

Ristra codes aim to be

- Asynchronous/Concurrent; Dynamic (compile- and run-time); Agile
- Open to new innovation in methods and CS

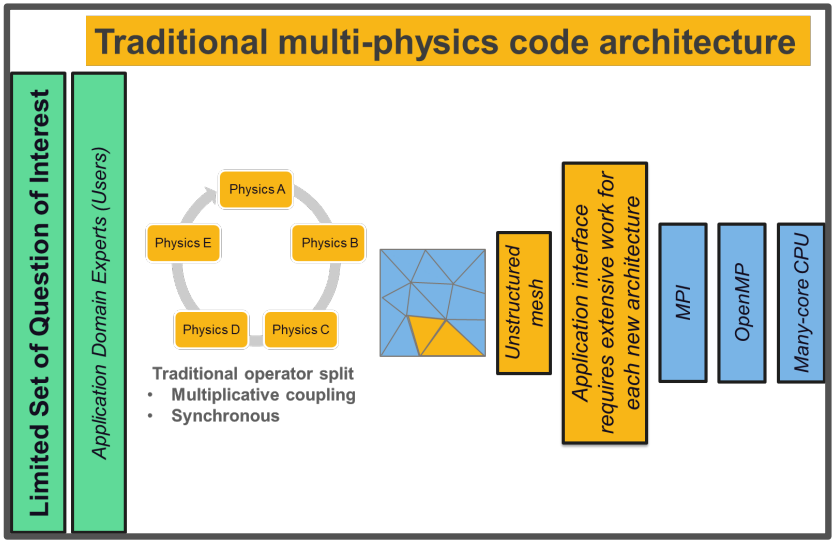

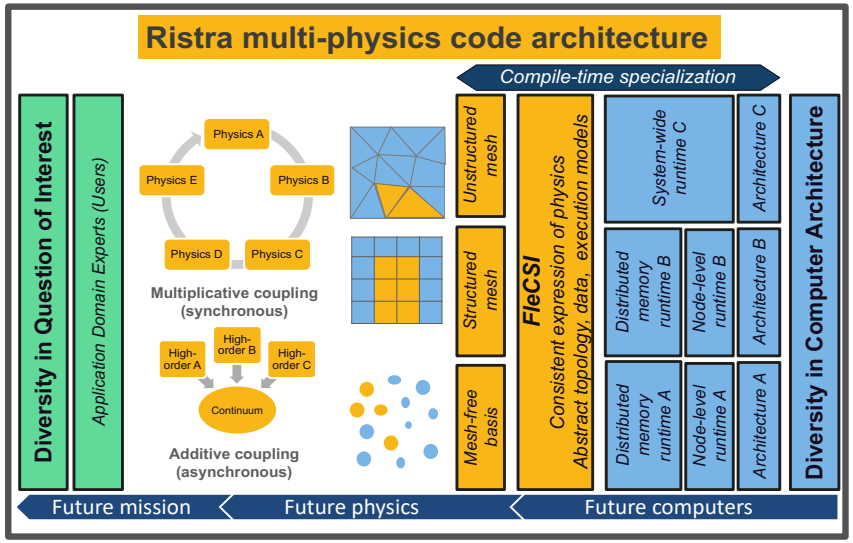

FleCSI: High-level Design

FleCSI is a C++ programming system for developing multiphysics simulation codes

- Runtime Abstraction Layer & Interfaces
  - High-level user interface → *application developers*
  - Mid-level static specialization interface → *computational scientists*
  - Core data structures → *computer scientists*
  - Support for multiple backend implementations
    → MPI, Legion, HPX (distributed-memory)
    → Kokkos, Kitsune (shared-memory)
- Programming Model (task-based, data-centric)
  - Runtime, control, execution, and data models
- Topologies
  - Mesh topologies (structured & unstructured)
  - Tree topologies
  - Set topologies

Control Model

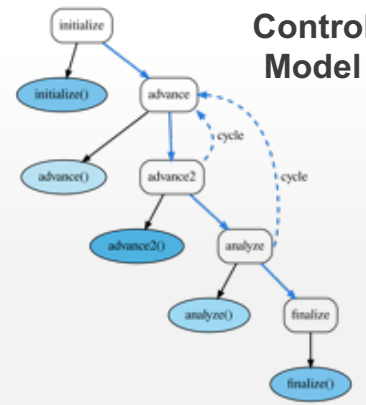

Separation of Concerns (Legion)

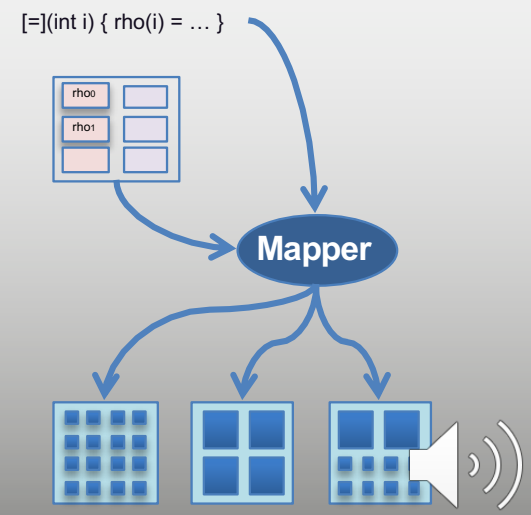

ATDM investments have pointed a way to improve efficiency in deploying codes on emerging architectures

- FleCSI abstraction layer helps automate porting to new platforms

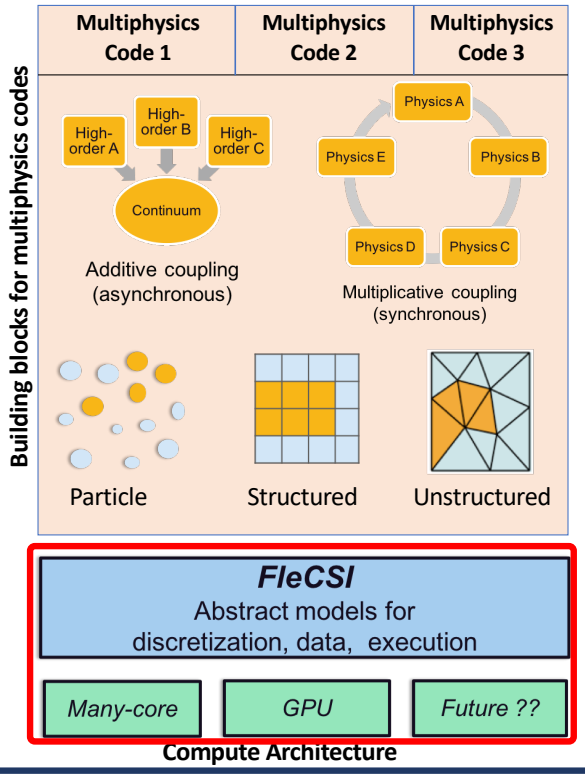

Each code ports to each architecture

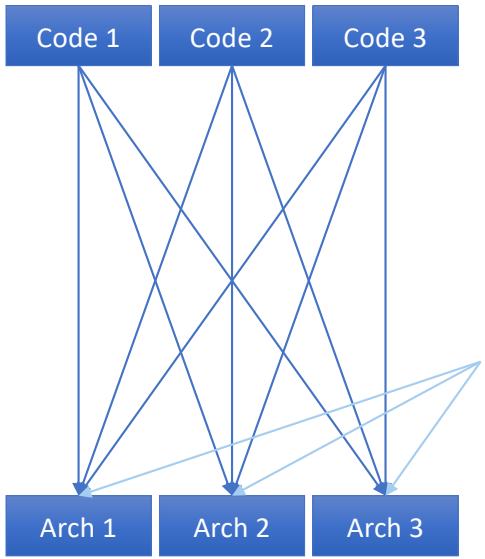

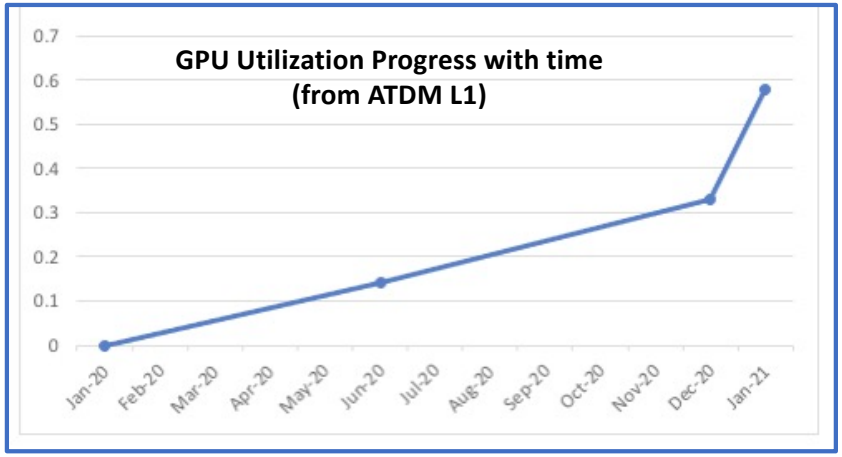

GPU Utilization in FleCSI-based code Puno on Sierra

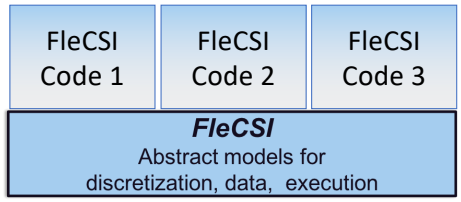

FleCSI's central model for data and execution can be ported once and all FleCSI-based codes benefit

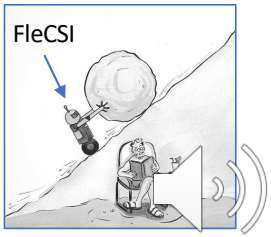

Through the ATDM program, we developed a powerful capability to explore new multiphysics methods

- ATDM investments also facilitate algorithmic exploration, in addition to improving efficiency of porting to emerging platforms

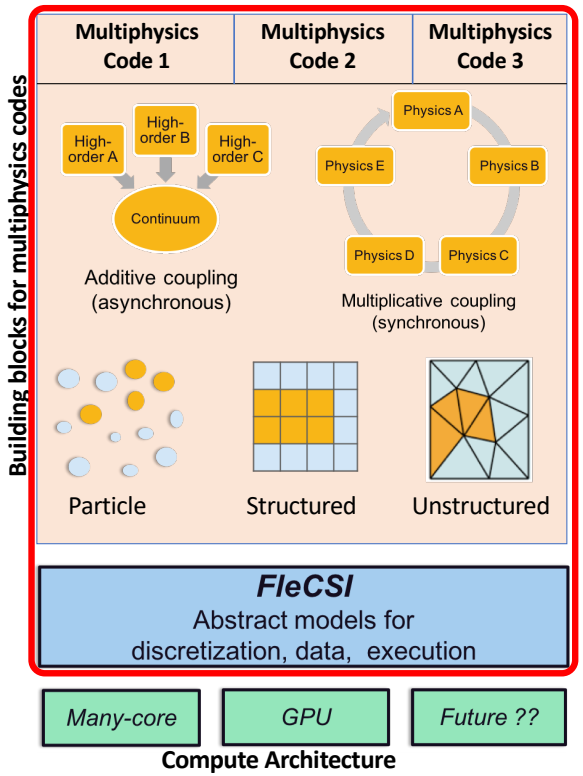

- disciplined and common development environment ease the <u>integration</u> of multiple physics package components
- Demonstration of multi-scale methods (HOLO Transport; VPSC material strength)
- Recent ATDM L1 demonstrated this agility for new physics algorithm development and deployment

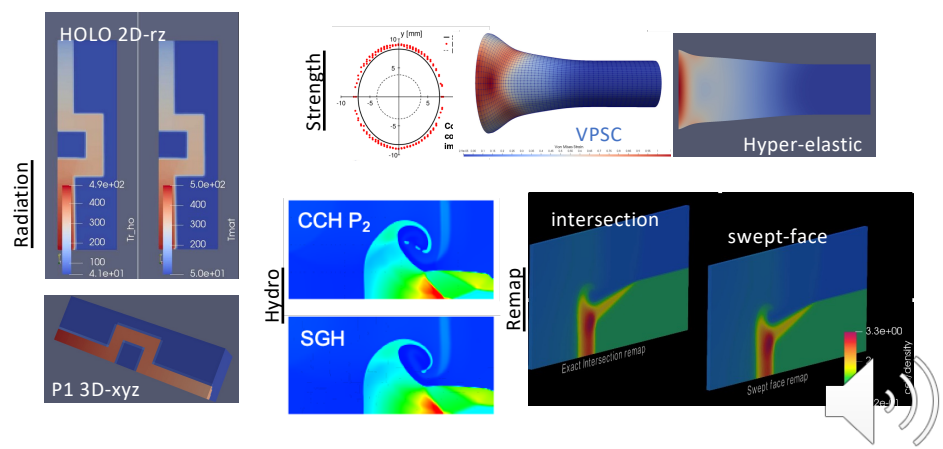

NGP: Our plan for advances in platform technology build on a recent industrial capability for specialization

The world is adopting customization as the path to performance in the post-Moore era

- ARM is at the center of a broad industrial capability for tailoring

Processor IP

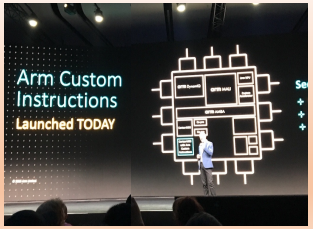

2.5D integration

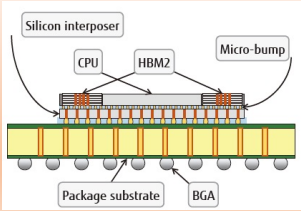

The Fab economy

GlobalShuttle MPW Schedule:

Tech. Node	Jan	Feb	Mar	Apr	May
12nm			18 MAR 1423		
22nm	13 JAN 2232		02 MAR 2233		11 MAY 2234

Japan and Europe have shown great potential for HPC

- US Exascale effort bet on a single path with its $2B investment

European Processor Initiative

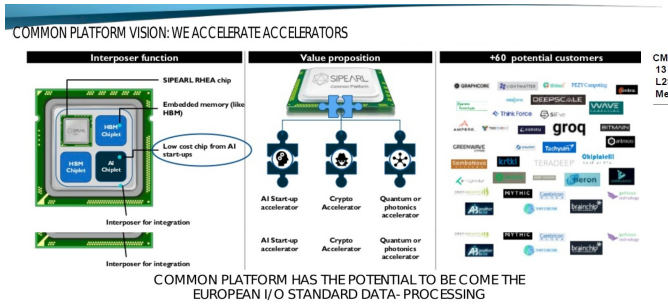

Japanese (MEXT) – Fugaku

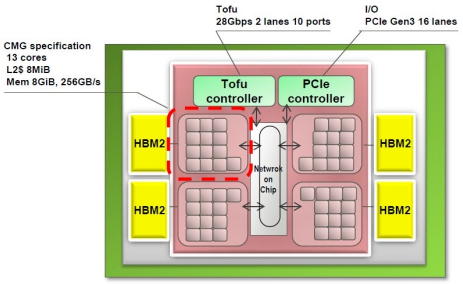

A new LANL institutional investment helps position the US on a path for long-term gains in computing

- Based on new line of NVIDIA technologies
  - $80M, delivery in early FY23
  - First new family in 2 decades
- Broad spectrum of advanced technologies for laboratory applications
  - Multiphysics, Machine Learning, Analytics…
  - Platform for hardware studies
- Since it is an institutional system, we can dedicate time to partnerships
  - Universities, DOD and industry

We and NVIDIA have assigned leaders for a codesign partnership

- Libraries and Compilers
- Deep Learning
- Storage
- Networking
- Light House Applications
- Sparsity, Simulations
- Characteristics + Profiling
- Architectures
- Testbeds

A partnership with NVIDIA answers the hardest question: How do we influence much larger industry efforts?

Summary

- Current mission drivers and machines are once again forcing exciting innovation in the computational tools we use for discovery science.
- LANL's recent ATDM efforts have laid the foundation for a code development environment that can meet those mission and machine drivers.
- Commercial trends are opening some great opportunities for transforming the nature of our computational toolkits
  - GPU machines are underwriting a boom in Machine Learning methodologies for a broad set of applications
  - ARM processors offer the opportunity for component tailoring efforts that is reinvigorating interest in deep codesign partnerships.