

# Overview of PSAAP III CS Areas/Exascale

**John Feddema**

**SNL Rep. to the ASC Alliance Strategy Team**

**PSAAP III Pre-proposal Conference**

**14-15 March 2018**


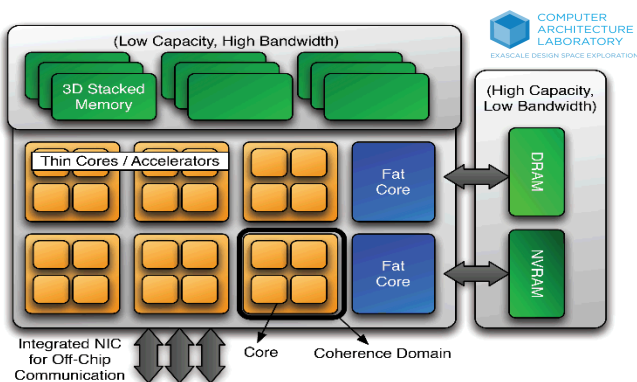
**Minneapolis, MN**

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525 SAND2018-XXXX PE



# Extreme-scale HPC architectures introduce programming challenges

System Change	Programming Challenge
Increased node-level parallelism	Expressing/managing node-level & hybrid parallelism
Diverse target architectures	Performance portability across systems
Decreased system reliability	Resilience/Fault mitigation
Increased system noise	Increased need for effective load-balancing strategies
Deeper memory hierarchies	Management of memory hierarchies/locality
Increased system scale	Increased workflow complexity




# **PSAAP III will support the following CS topics and more:**

- **Exploration of advanced HPC architectures**
- **Programming environments and runtime systems**
- **Productivity & Performance Portability**
- **Workflow automation**
- **Data Analytics for science and engineering applications**
- **New approaches to engineering**
- **Algorithms/models**

# Exploration of Advanced HPC Architectures


- **Architecture simulation/emul**
  - Multi-level memory managemen
  - CPU/GPU/Neuromorphic/...
  - MPI, Networking
- **Performance prediction**

AMD Two Level Memory Concept

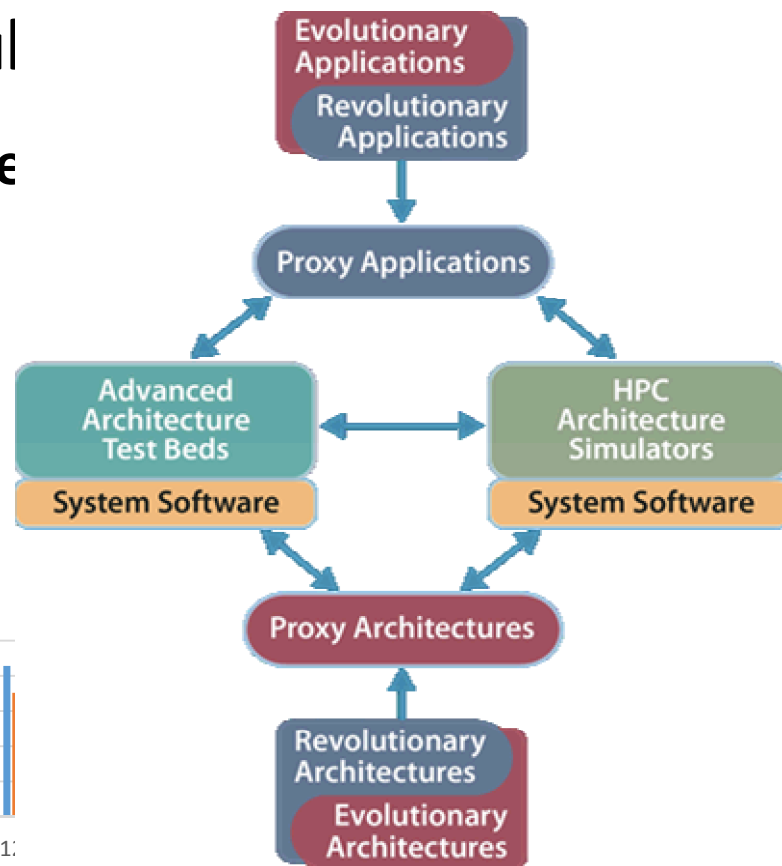


Two types of memory (high-bandwidth for fast access, low-cost for capacity)

miniFE

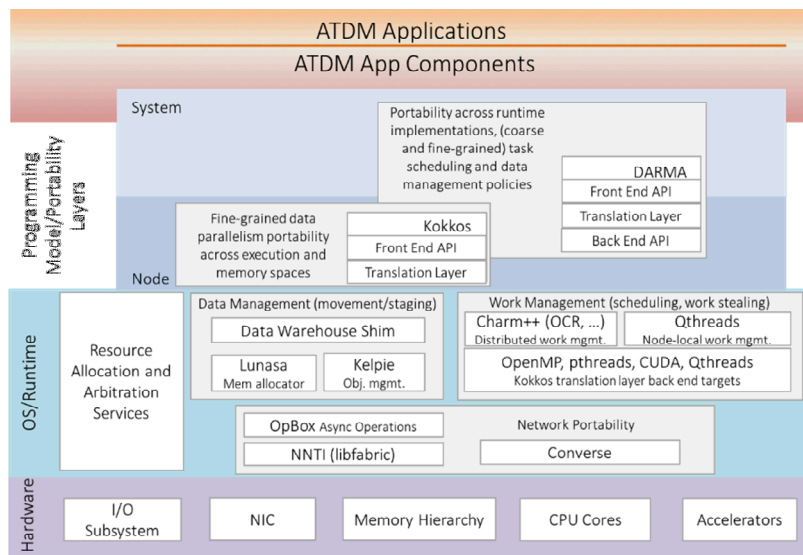


A mixed controller (in blue) is able to match or outperform a single controller on proxy scientific application program flow.



# Programming Environments and Runtime Systems

- Composition of libraries, runtimes, programming languages
- Performance portability layers to insulate users from hardware





## SNL-Lead Kokkos and DARMA efforts

\*ATDM = Advanced Technology Demonstration and Mitigation

\*\*DARMA = Distributed Asynchronous Resilient Models and Applications

*Other University-led and National Lab-Led efforts*


HPX-5 


Habanero-UPC++ 

OCR  

Legion   

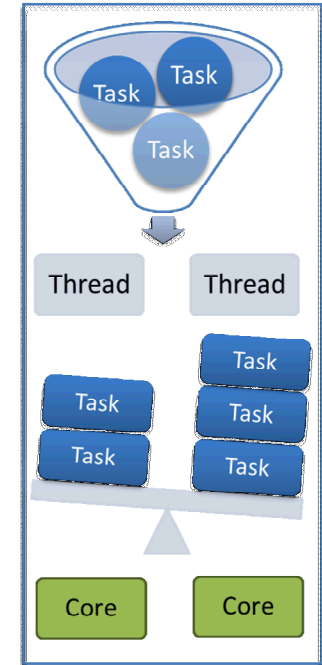
Uintah 

CHAI+RAJA 

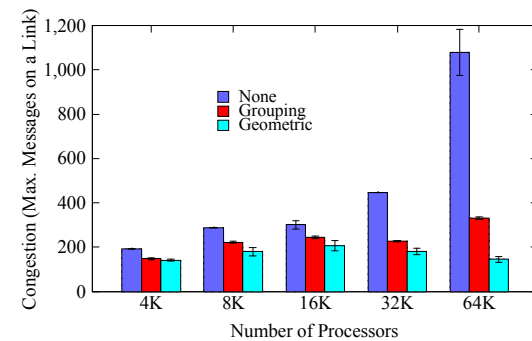


Charm++ 

# Programming Environments and Runtime Systems (Continued)

- Task-based programming
  - Interoperability of Asynchronous Many Task (AMT) with MPI and Threading
  - Underlying mechanisms to support load balancing in AMT



- MPI Rank Placement

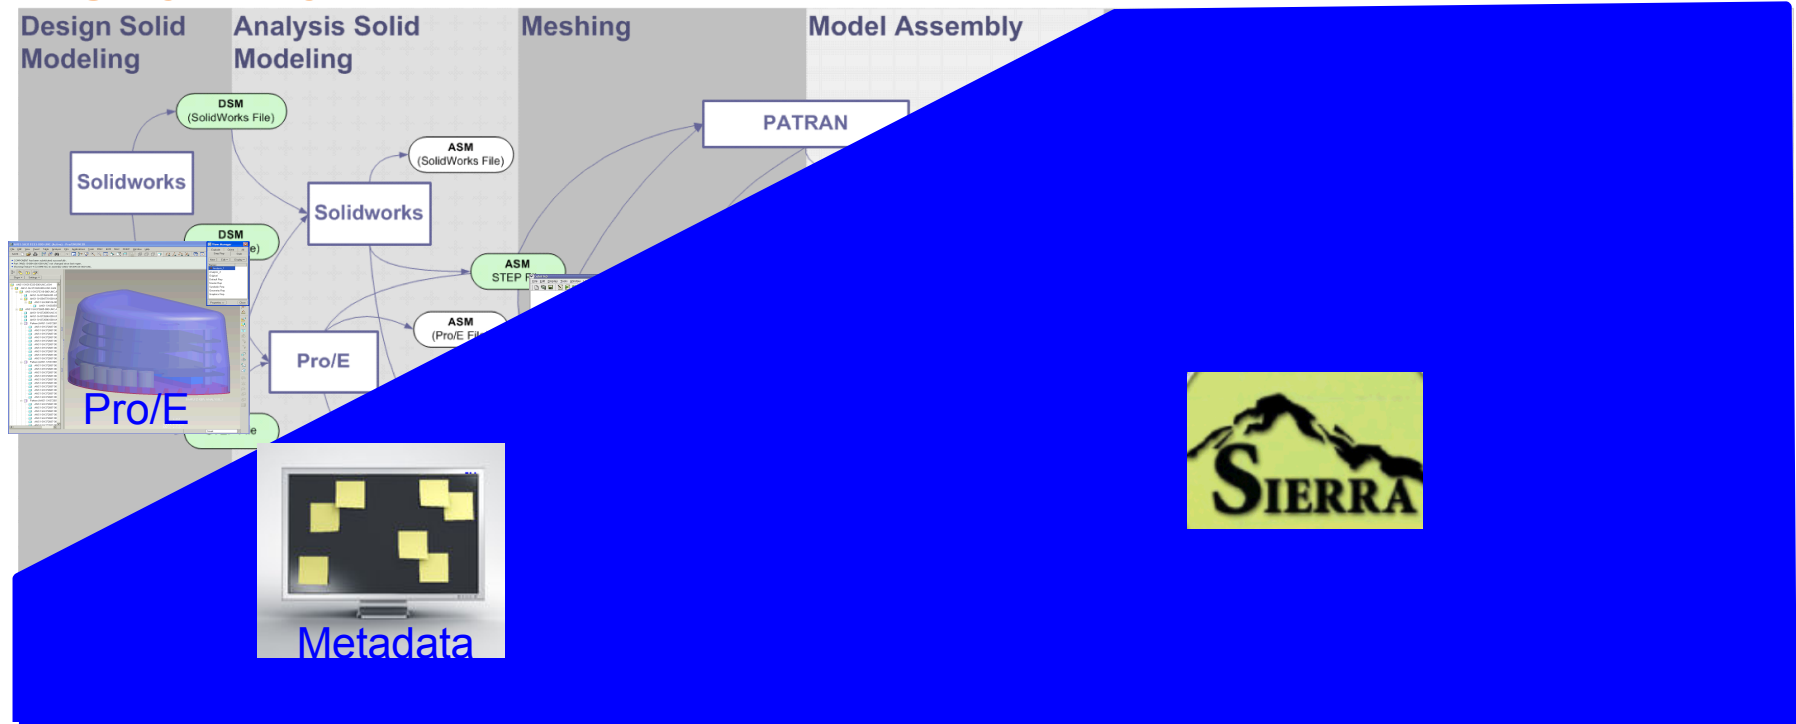


Geometric task placement decreases the distance messages travel and the amount of congestion in the network. Ref: Devine

# Analysis Problem – Many Tools and Many Files

## Legacy Analysis Process/Data Workflow

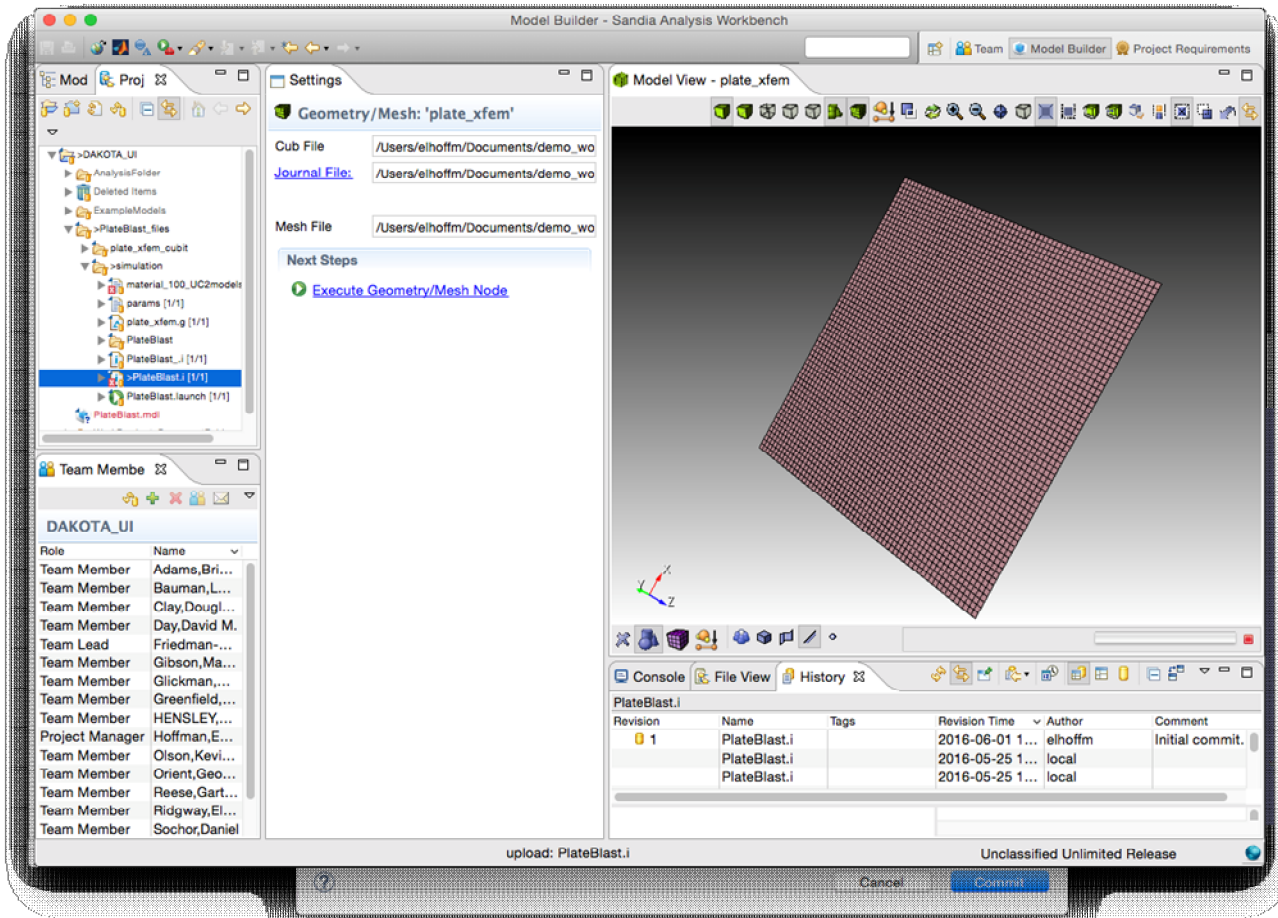
## Sandia Analysis Workbench




- Involved many (stove-piped) tools → **High training costs**
- Large number of artifacts, managed by the analyst → **provenance and process knowledge is lost**
- Models are continuously getting more complex with larger datasets

# Sandia Analysis Workbench (SAW) Streamlines the Analysis Process while Retaining Analysis Data

Multiple tools and services integrated in an Eclipse\* workbench application



**Model Building**  
**Process**  
**Management**  
**Data & Metadata**  
**Management**




\* <http://www.eclipse.org>

# Productivity & Performance Portability

- **Performance Portable Programming models**
  - Abstractions such as Kokkos and RAJA
- **Environments for efficient development of simulations**
  - e.g., Integrating tools and services into Eclipse IDE
- **Techniques for adding new physics**
  - e.g., Drekar: Resistive MHD, multifluid EM plasmas with coupled multi-physics
- **Hiding the complexity of heterogeneous computing systems**

# Workflow Automation

- Simulation setup, simulation runs or complex ensemble runs and post processing
- Management of bulk data
  - Improve data management service
  - Use nonvolatile memory to speed up data handoffs
- Simulation repeatability, e.g., role of containers
- Dynamic resource management, e.g., UQ pipeline




Advanced Technology Development and Mitigation (ATDM) Data Warehouse project  
Ref: Craig Ulmer, Ron Oldfield



# Data Analytics for Science/Engineering Applications (continued)


- **Machine learning (ML)**
  - Quantifying uncertainty in ML
  - Neuromorphic computing
  - Impact of embedded ML on application performance
- **Rigorous math models for data analytics**




ML Models of Turbulent Flow UQ  
Ref: Ling, Barone, Chowdhary, Fike, Davis, et. al.

# New Approaches to Engineering

- Machine learning applied to design
  - E.g., topological design optimization of Z-pinch convolute



Goal: Optimal design of Z-pinch Convolute



Iterations in the Design of a Coaxial Cable  
Ref: Cyr, von Winkel, Kouri, Ridzal, Shadid, Miller, 2017



- Models and processes for the physics and chemistry of additive manufacturing

# Large Spatial Resolution Motivates Multi-scale/physics Simulation with Optimization under Uncertainty Capabilities

**Challenges:**  
 multiscale  
 multiphysics  
 large optimization space  
 manage uncertainties


Part Scale Thermal & Solid Mechanics

Mesoscale Thermal Behavior




Part Scale Microstructure



Powder Spreading



Powder Behavior



Mesoscale Texture/Solid Mechanics/CX



$10^{-6}$

$10^{-3}$   
 Length Scale (m)

1

Ref: Bart van Bloemen Waanders,  
 Timothy Wildey, Daniel T. Seidl

# Algorithms and Models

- **Novel approaches to coupling multiphysics/multiscale**
- **Algorithms for increasing performance of HPC systems, e.g. latency hiding, reduction of synchronization, utilization of simultaneous execution**
- **Support for resilience**
- **Exposing more parallelism at the cost of algorithm efficiency**
- **Reduced order models**
- **Stochastic algorithms and adaptive algorithms**

# CS/Exascale Components in Multidisciplinary Centers and Single-Discipline Centers

**MSC**

Overarching application

Science /engineering disciplines

V&V/UQ

Exascale CS

**SDC**

Integrated problem/challenge

Research directions within single discipline

V&V/UQ

Exascale CS

or

**SDC**

Exascale CS