



# The Rio Grande Consortium for Advanced Research on Exascale Simulation (Grande CARES)

Dr. Vinod Kumar, Professor and Department Chair  
Mechanical and Industrial Engineering  
Texas A&M University, Kingsville, TX

Aug. 22, 2024

Annual Symposium: AI Accelerated Physics Based Modelling  
and its Role in Energy Industry

Hosted by: Computational Fluid Dynamics  
SLB Q-Auditorium, 10001 Richmond Ave, Houston, TX







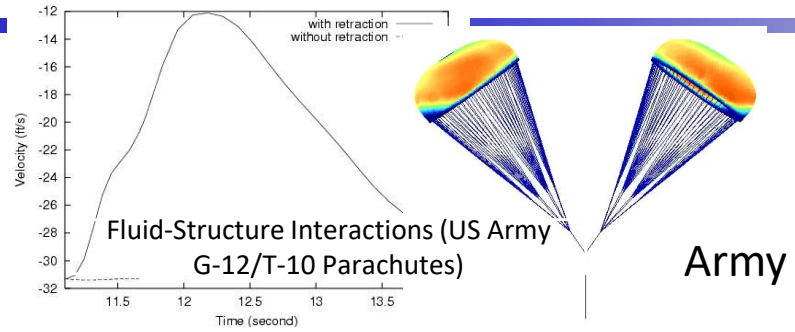
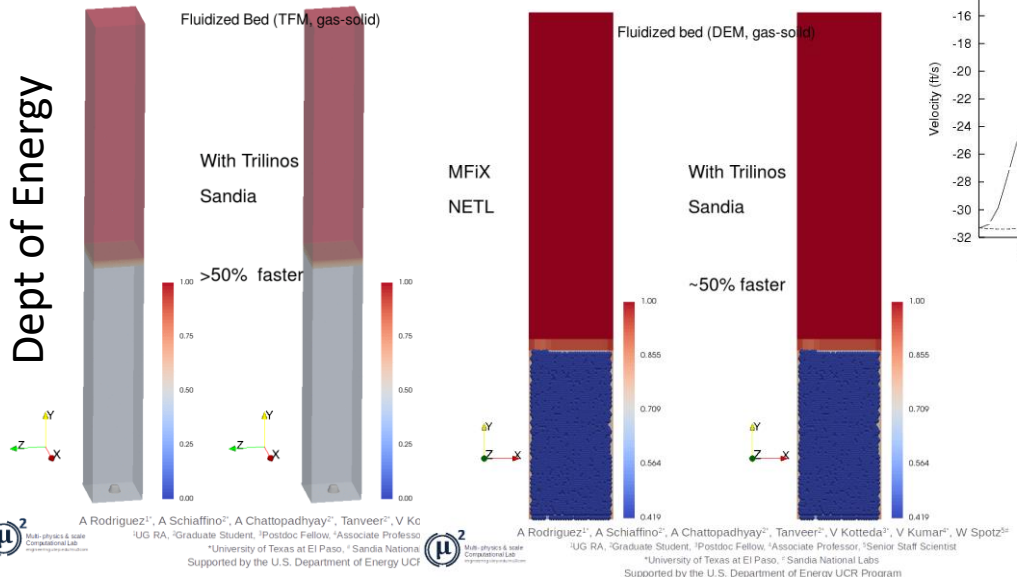
# Outline

---



- **Background**
- **Methodology and Technical Discussion**
- **Crosscutting themes**
  - **HPC & Exascale computing – Sandia National Labs & our software development activities**
  - **Machine Learning, AI, and UQ**
- **Selected Problems**
- **Concluding remarks**

# Selected Projects

Dept of Energy


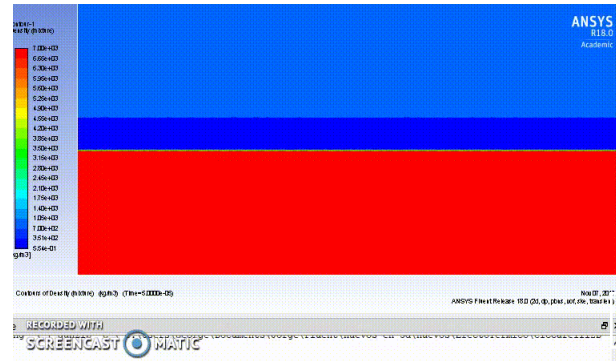





**Uncertainty Quantification**  
**Materials for Extreme Environments (AFOSR)**  
 Exa-scale Pore Network Simulator (EXPNS):  
 High viscosity fluid invading through porous media




A Rodriguez<sup>1</sup>, A Chattopadhyay<sup>2</sup>, Tanveer<sup>2</sup>, V Kottedda<sup>3</sup>, V Kumar<sup>4</sup>  
<sup>1</sup>Graduate Student, <sup>2</sup>Postdoc Fellow, <sup>3</sup>Associate Professor, <sup>4</sup>Senior Staff Scientist  
<sup>1</sup>University of Texas at El Paso, <sup>2</sup>Sandia National Labs  
 Supported by the U.S. Department of Energy UCR Program

Air Force



Air Force



# Compressible Flows in Engineer's Language (Math)

PDEs, ODEs are obsolete! Easier than learning Spanish!

On  $\Omega_t \forall t \in (0, T)$

Density

Velocity vector

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0$$

Pressure

Stress-tensor

Conservation of mass/Continuity equations

$$\frac{\partial(\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \mathbf{u}) - \nabla p - \nabla \cdot \mathbf{T} = \mathbf{0}$$

Momentum conservation

$$\frac{\partial(\rho e)}{\partial t} + \nabla \cdot (\rho e \mathbf{u}) - \nabla \cdot (\mathbf{T} \mathbf{u}) - \nabla \cdot \mathbf{q} = 0$$

Energy conservation

Energy per unit volume

Heat-flux

Internal energy per unit volume

$$\begin{bmatrix} T_{11} & T_{12} & T_{13} \\ T_{21} & T_{22} & T_{23} \\ T_{31} & T_{32} & T_{33} \end{bmatrix} = \mathbf{T} = \mu \left( (\nabla \mathbf{u}) + (\nabla \mathbf{u})^T \right) - \frac{2}{3} \mu (\nabla \cdot \mathbf{u}) \mathbf{I}$$

Constitutive relationship

Viscosity

Ideal gas equations

$$p = (\gamma - 1) \rho i = \rho R \theta,$$

$$i = e - \frac{1}{2} \|\mathbf{u}\|^2 = C_v \theta$$

Specific gas constant

$$C_v = \frac{R}{\gamma - 1}, \quad C_p = \frac{\gamma R}{\gamma - 1},$$

Heat Conduction

Energy  $\rightarrow$  Heat Equation

$$\mathbf{q} = -\kappa \nabla \theta \Rightarrow \nabla \cdot \kappa \nabla \theta = 0$$

Conductivity

Temperature

$$\text{Reynolds number, } Re = \frac{\rho_o U D}{\mu}, \quad \text{Prandtl number, } Pr = \frac{\mu C_p}{\kappa}$$

# Mesh and Membrane

## Mesh Motion

$$\nabla \cdot \boldsymbol{\sigma}^m + \mathbf{f}^m = \mathbf{0} \text{ on } \Omega_t^m$$

$$\boldsymbol{\sigma}^m = \lambda^m \text{tr}(\boldsymbol{\varepsilon}^m) \mathbf{I} + 2\mu^m \boldsymbol{\varepsilon}^m,$$

$$\boldsymbol{\varepsilon}^m = \frac{1}{2} \left( (\nabla \mathbf{y}^m) + (\nabla \mathbf{y}^m)^T \right)$$

$$\mathbf{y}^m = \mathbf{g}^m \text{ on } (\Gamma_t^m)_g,$$

$$\mathbf{n} \cdot \boldsymbol{\sigma}^m = \mathbf{h}^m \text{ on } (\Gamma_t^m)_h$$

Equations of motions: Pseudo-solid

Stress-strain relationship:

Linear elastic solid

Boundary conditions

## Structural Dynamics

$$\rho^s \left( \frac{d^2 \mathbf{y}^s}{dt^2} + \eta \frac{d\mathbf{y}^s}{dt} - \mathbf{f}^s \right) - \nabla \cdot \boldsymbol{\sigma}^s = \mathbf{0} \text{ on } \Omega_t^s$$

← Force balance

Piola-Kirchoff Stress:  $\mathbf{S} = \frac{\rho_0^s}{\rho_t^s} \mathbf{F} \boldsymbol{\sigma}^s \mathbf{F}^T$ , Deformation Gradient:  $F_{ij} = \mathbf{g}_{ik} G^{kj}$

Constitutive relationship:  $S^{ij} = \left( \frac{2\lambda^s \mu^s}{\lambda^s + 2\mu^s} G^{ij} G^{kl} + \mu^s [G^{il} G^{jk} + G^{ik} G^{jl}] \right) \varepsilon_{kl}$

BCs:  $\mathbf{y}^s = \mathbf{g}^s$  (Dirichlet),  $\mathbf{n} \cdot \boldsymbol{\sigma}^s = \mathbf{h}^s$  (Neumann's),

Initial Conditions:  $\mathbf{y}^s = \mathbf{0} \equiv \frac{d\mathbf{y}^s}{dt}$  on  $(\Omega_0^s)_g$

$\boldsymbol{\sigma}$ : Cauchy stress

$\mathbf{g}_{ij}$ : Covariant metric in deformed configuration

$G_{ij}$ : Contravariant metric in original configuration

Lame's constant:  $\lambda^s, \mu^s$



# SST/DSD Formulations

$$\begin{aligned} & \int_{Q_n} \mathbf{W}^h \cdot \left( \frac{\partial \mathbf{U}^h}{\partial t} + \mathbf{A}_i^h \cdot \frac{\partial \mathbf{U}^h}{\partial x_i} \right) dQ + \int_{Q_n} \left( \frac{\partial \mathbf{W}^h}{\partial x_i} \right) \cdot \left( \mathbf{K}_{ij}^h \frac{\partial \mathbf{U}^h}{\partial x_j} \right) dQ \\ & + \int_{\Omega_n} (\mathbf{W}^h)_n^+ \cdot \rho((\mathbf{U}^h)_n^+ - (\mathbf{U}^h)_n^-) dQ \\ & + \sum_{e=1}^{(n_{el})_n} \int_{Q_n^e} \tau_{\text{SUPG}} \left[ \frac{\partial \mathbf{W}^h}{\partial t} + (\mathbf{A}_k^h)^T \frac{\partial \mathbf{W}^h}{\partial x_k} - \frac{\partial}{\partial x_k} \left( \mathbf{K}_{kl}^h \frac{\partial \mathbf{W}^h}{\partial x_l} \right) \right] \cdot \\ & \left[ \frac{\partial \mathbf{U}^h}{\partial t} + (\mathbf{A}_i^h) \frac{\partial \mathbf{U}^h}{\partial x_i} - \frac{\partial}{\partial x_i} \left( \mathbf{K}_{ij}^h \frac{\partial \mathbf{U}^h}{\partial x_j} \right) \right] dQ \\ & + \sum_{e=1}^{(n_{el})_n} \int_{Q_n^e} \tau_{\text{SHOC}} \left( \frac{\partial \mathbf{W}^h}{\partial x_i} \right) \cdot \left( \frac{\partial \mathbf{U}^h}{\partial x_i} \right) dQ = \int_{(P_n)_h} \mathbf{W}^h \cdot \mathbf{h}^h dP. \end{aligned}$$

# Multiphysics Framework – Strong coupling

Couple fluid dynamics with physics from various disciplines for realism, e.g., **Fluid Dynamics**, Structural Dynamics, Rain/Particle Dynamics, Solid Mechanics, Meteorology, Oceanography

$$\rho^s \left( \frac{d^2 \mathbf{y}^s}{dt^2} + \eta \frac{d\mathbf{y}^s}{dt} - \mathbf{f}^s \right) - \nabla \cdot \boldsymbol{\sigma}^s = \mathbf{0} \quad \text{Physics 1 (P}_1\text{)}$$

$$\nabla \cdot \boldsymbol{\sigma}^m + \mathbf{f}^m = \mathbf{0} \quad \text{Physics 2 (P}_2\text{)}$$

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \quad \text{Physics 3 (P}_3\text{)}$$

$$\frac{\partial (\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \mathbf{u}) - \nabla p - \nabla \cdot \mathbf{T} = \mathbf{0}$$

After discretizing using your favorite method (FVM / FEM)

$$\begin{aligned} N_1(\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3) &= 0 \\ N_2(\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3) &= 0 \\ N_3(\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3) &= 0 \end{aligned}$$

Linearization of nonlinear equations – Newton Raphson

$$\begin{aligned} \left[ \frac{\partial N_1}{\partial \mathbf{d}_1} \right]^i \Delta \mathbf{d}_1^{i+1} + \left[ \frac{\partial N_1}{\partial \mathbf{d}_2} \right]^i \Delta \mathbf{d}_2^{i+1} + \left[ \frac{\partial N_1}{\partial \mathbf{d}_3} \right]^i \Delta \mathbf{d}_3^{i+1} &= -[N_1(\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3)]^i \\ \left[ \frac{\partial N_2}{\partial \mathbf{d}_1} \right]^i \Delta \mathbf{d}_1^{i+1} + \left[ \frac{\partial N_2}{\partial \mathbf{d}_2} \right]^i \Delta \mathbf{d}_2^{i+1} + \left[ \frac{\partial N_2}{\partial \mathbf{d}_3} \right]^i \Delta \mathbf{d}_3^{i+1} &= -[N_2(\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3)]^i \\ \left[ \frac{\partial N_3}{\partial \mathbf{d}_1} \right]^i \Delta \mathbf{d}_1^{i+1} + \left[ \frac{\partial N_3}{\partial \mathbf{d}_2} \right]^i \Delta \mathbf{d}_2^{i+1} + \left[ \frac{\partial N_3}{\partial \mathbf{d}_3} \right]^i \Delta \mathbf{d}_3^{i+1} &= -[N_3(\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3)]^i \end{aligned}$$

After linearization (two way coupling)

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} & \mathbf{A}_{13} \\ \mathbf{A}_{21} & \mathbf{A}_{22} & \mathbf{A}_{23} \\ \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \end{bmatrix}^i \begin{bmatrix} \Delta \mathbf{d}_1 \\ \Delta \mathbf{d}_2 \\ \Delta \mathbf{d}_3 \end{bmatrix}^{i+1} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{b}_3 \end{bmatrix}^i$$

One way coupling – used in my FSI & most multiphysics framework (e.g., ANSYS, COMSOL) **Not good for tight coupling cases**

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} & \mathbf{0} \\ \mathbf{A}_{21} & \mathbf{A}_{22} & \mathbf{0} \\ \mathbf{A}_{31} & \mathbf{A}_{32} & \mathbf{A}_{33} \end{bmatrix}^i \begin{bmatrix} \Delta \mathbf{d}_1 \\ \Delta \mathbf{d}_2 \\ \Delta \mathbf{d}_3 \end{bmatrix}^{i+1} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{b}_3 \end{bmatrix}^i$$

Generalization: Let Physics1, Physics2, Physics3, ..., Physics-n be represented by P<sub>1</sub>, P<sub>2</sub>, P<sub>3</sub>, ..., P<sub>n</sub> respectively.


$$\begin{bmatrix} \mathbf{A}_{P_1 P_1} & & & & \\ \mathbf{A}_{P_2 P_1} & \mathbf{A}_{P_2 P_2} & & & \\ \cdot & \cdot & \cdot & & \\ \cdot & \cdot & \cdot & \cdot & \\ \mathbf{A}_{P_n P_1} & \mathbf{A}_{P_n P_2} & \cdot & \cdot & \mathbf{A}_{P_n P_n} \end{bmatrix} \begin{bmatrix} \Delta \mathbf{d}_{P_1} \\ \Delta \mathbf{d}_{P_2} \\ \cdot \\ \cdot \\ \Delta \mathbf{d}_{P_n} \end{bmatrix} = \begin{bmatrix} \Delta \mathbf{b}_{P_1} \\ \Delta \mathbf{b}_{P_2} \\ \cdot \\ \cdot \\ \Delta \mathbf{b}_{P_n} \end{bmatrix}$$

# DOE/NNSA Exascale Project

<https://sites.google.com/view/grande-cares/home>

## The Rio Grande Consortium for Advanced Research on Exascale Simulation (Grande CARES)

\$5M, 2022-2027, DOE NNSA/MSIPP (Grant Number: GRANT13584020)



Understanding complex physics and physical phenomena of intricate engineering systems is critical for designing and maintaining reliable, efficient, and economic systems that can be safely operated. System-level engineering challenges are often complex and involve multi-physics coupling of multiple scales in several orders of magnitude. The latest advancement in computational and data-driven technologies with supercomputing interfaces/infrastructures promises novel ways to address intricate engineering challenges. In addition, the advanced modeling & simulation (M&S) capabilities allow in-depth investigations and can significantly reduce operational costs. However, effective implementation of these technologies requires a workforce educated with an in-depth understanding of multi-physics concepts from multiple disciplines & cross-cutting technologies. The CARES team aims to address this gap through an innovative multi-physics integrator comprising five core research thrusts and an innovative curriculum. The educational goal is to cultivate scientists and engineers from underrepresented groups, educated in integrative skills needed for the advancement of M&S. The research will focus on developing and integrating cutting-edge computational algorithms using HPC, ML, data analytics, UQ, and other novel computational capabilities. The consortium brings together researchers from 5 departments (Mechanical Engineering, Civil Engineering, Computer Science, Computational Science, & Mathematical Science). It is comprised of faculty and staff from 4 prominent HSI (UTEP, UNM, NMSU, and NMT), one HBCU (PVAMU), and Sandia (a prominent DOE NNSA lab with leading-edge advanced M&S capabilities).



# CARES: Education

- **Interdisciplinary Education / Curriculum**

*Courses:* Computational Methods for Multiphysics Problems, Advanced Computing, Data-driven analytics and engineering


*Internship experience:*

*Interdisciplinary project*

- **Interdisciplinary Student Research**

Research leadership, Written communication and accountability, Oral communication, Advanced computing workshops, Seminar

- **Mentoring**



# CARES: Research

- Formulate, validate and implement advanced computational tools operating efficiently with extremely large databases for system level predictions.
- Develop high-fidelity, multiphysics computational approaches for coupled thermal-fluid-structural systems incorporating computational fluid & structural mechanics and high temperature reactivity.
- Deploy leading-edge computational tools from SNL through integrative curriculum changes to develop a sustainable approach to workforce needs trained in advanced M&S.

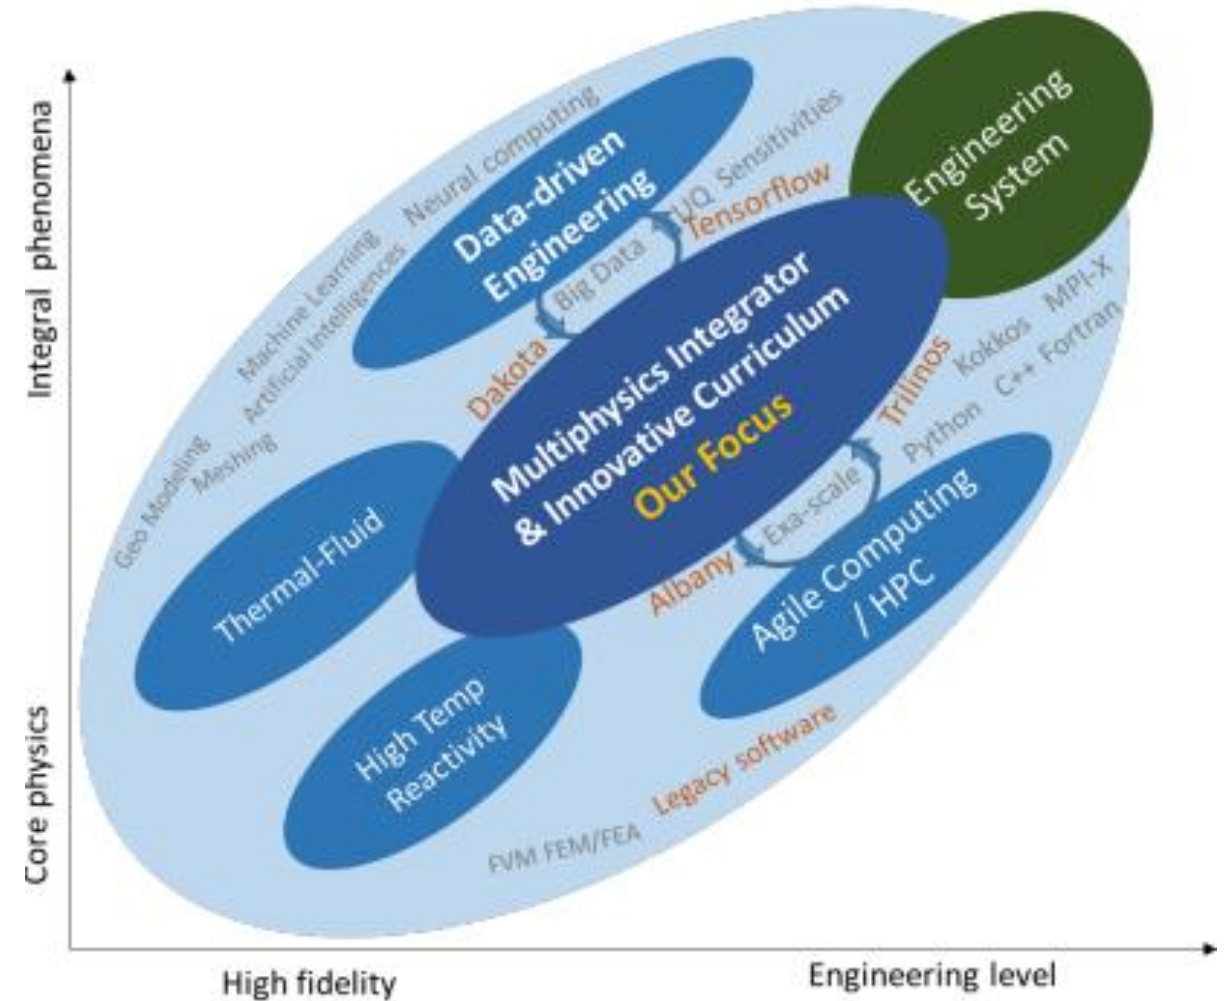
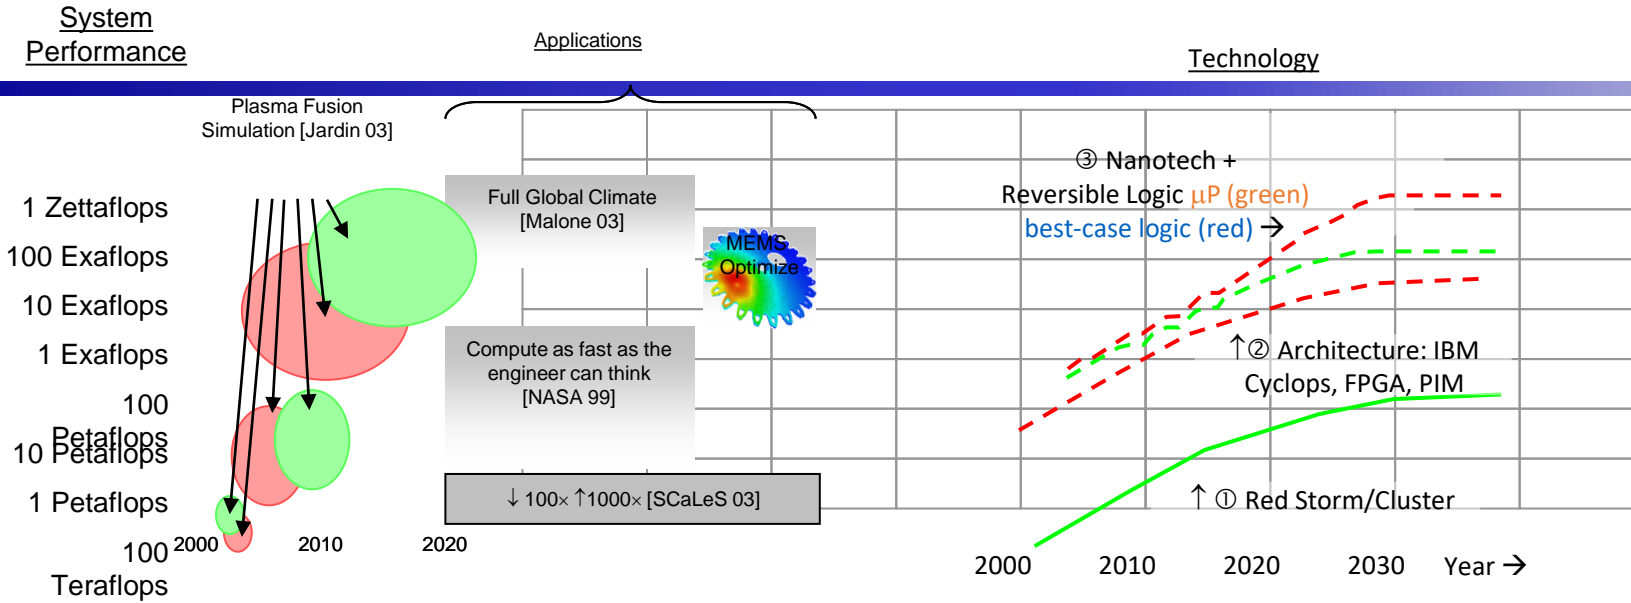


Figure 1. Rio Grande CARES research focus and thrusts.

Today's HPC infrastructure are a lot more complex but very affordable – GPU, Threading, MPI-X  
Thanks to thriving video gaming industries!

Slide Credits: DeBenedictis, Erik P, "Petaflops Exaflops and Zettaflops for Climate Modeling."



## Onyx Cray

Login cores: 12 x 22 node/core

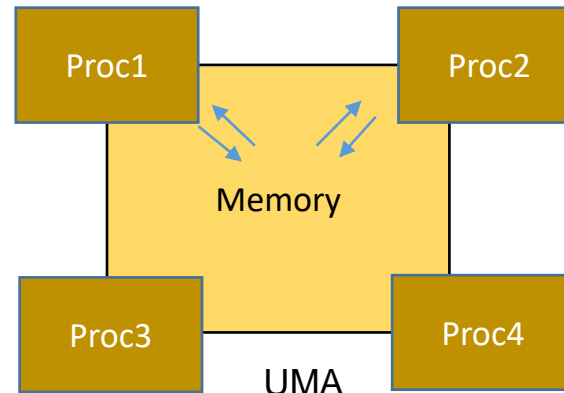
Standard cores: 2,858 x 44 node/core

Memory: 122/998Gbytes/core

Accelerators: KNL & GPU




### Shared Memory Model



$$S_p = \frac{T_p}{T_s} = \frac{1}{F_s + \frac{F_p}{P} + T_c}$$


### Distributed Memory Model



NUMA

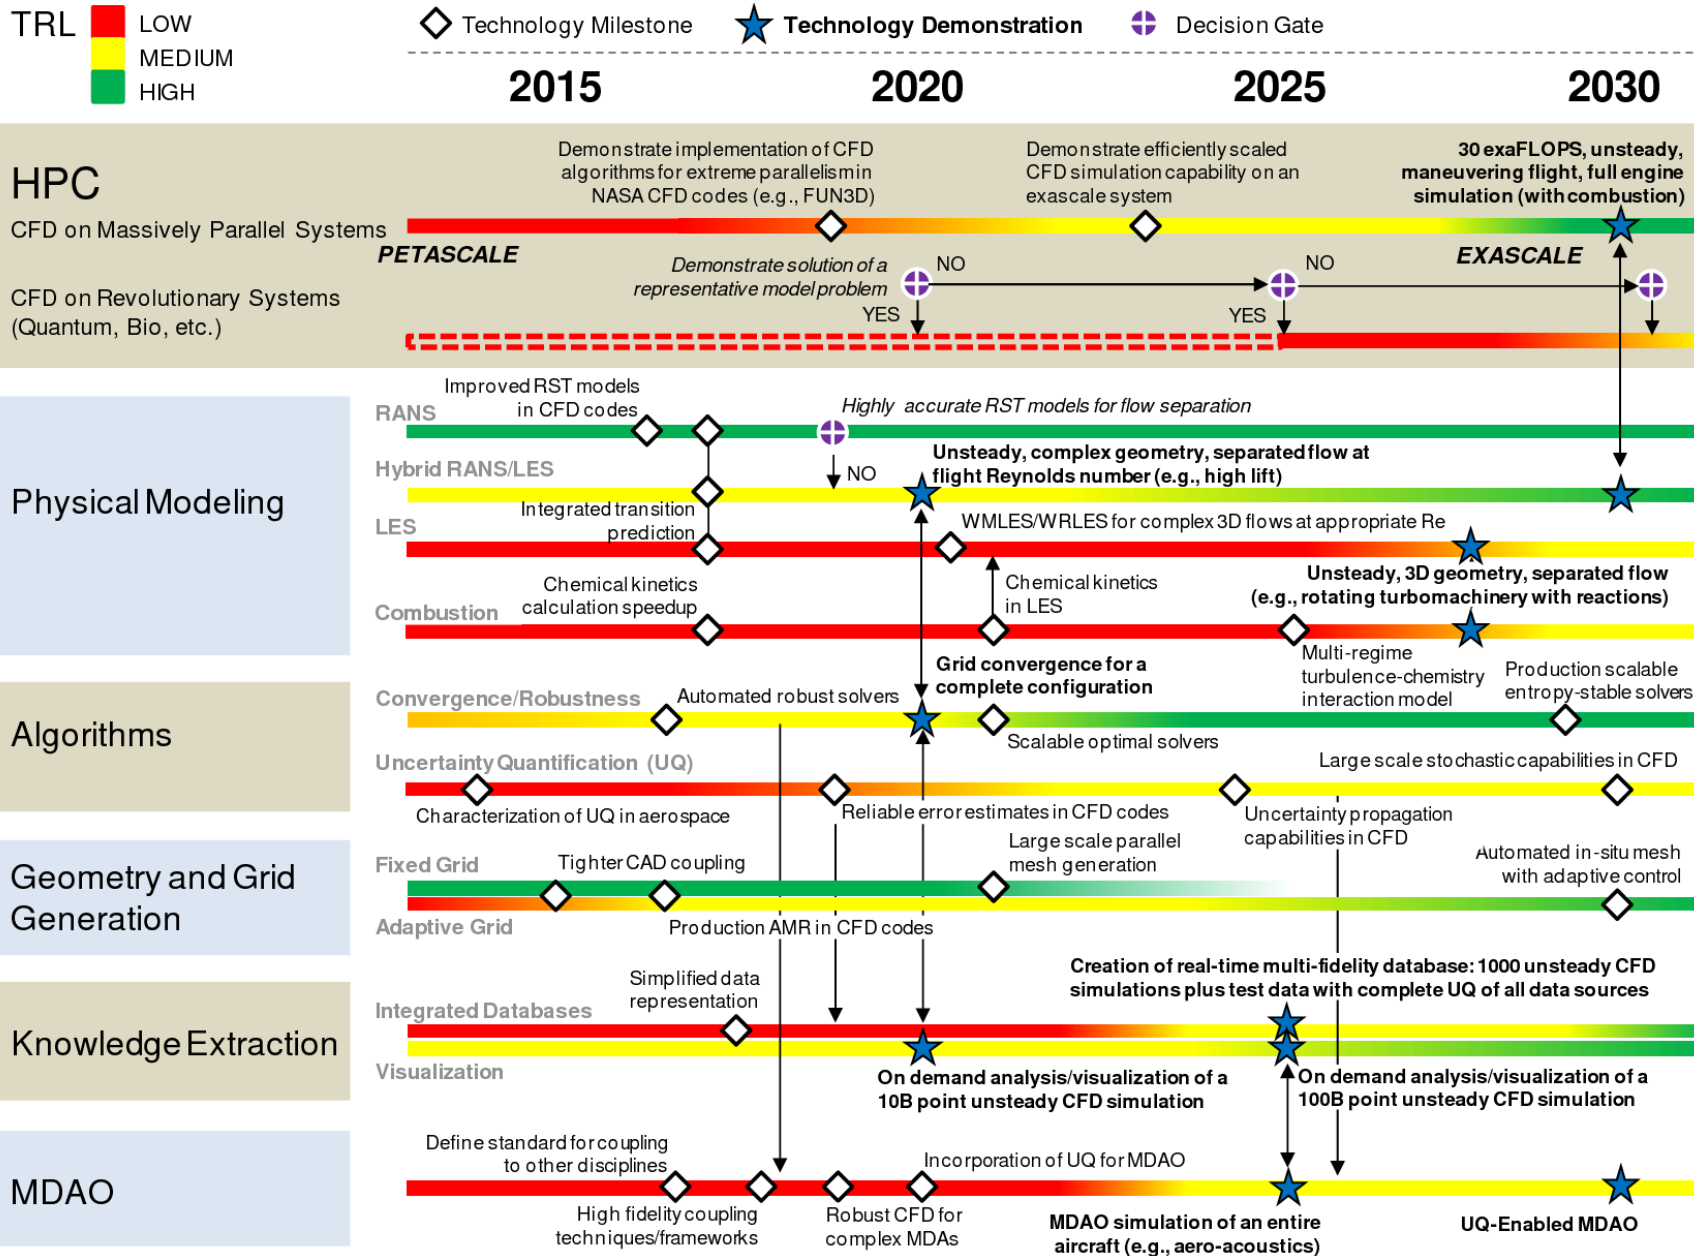
### Accelerators

- Light weight computing
- Thousands of threads
- Huge speed up
- Host-Device transfer is the bottleneck



# CFD Vision 2030 Study : A Path to Revolutionary Computational Aerosciences

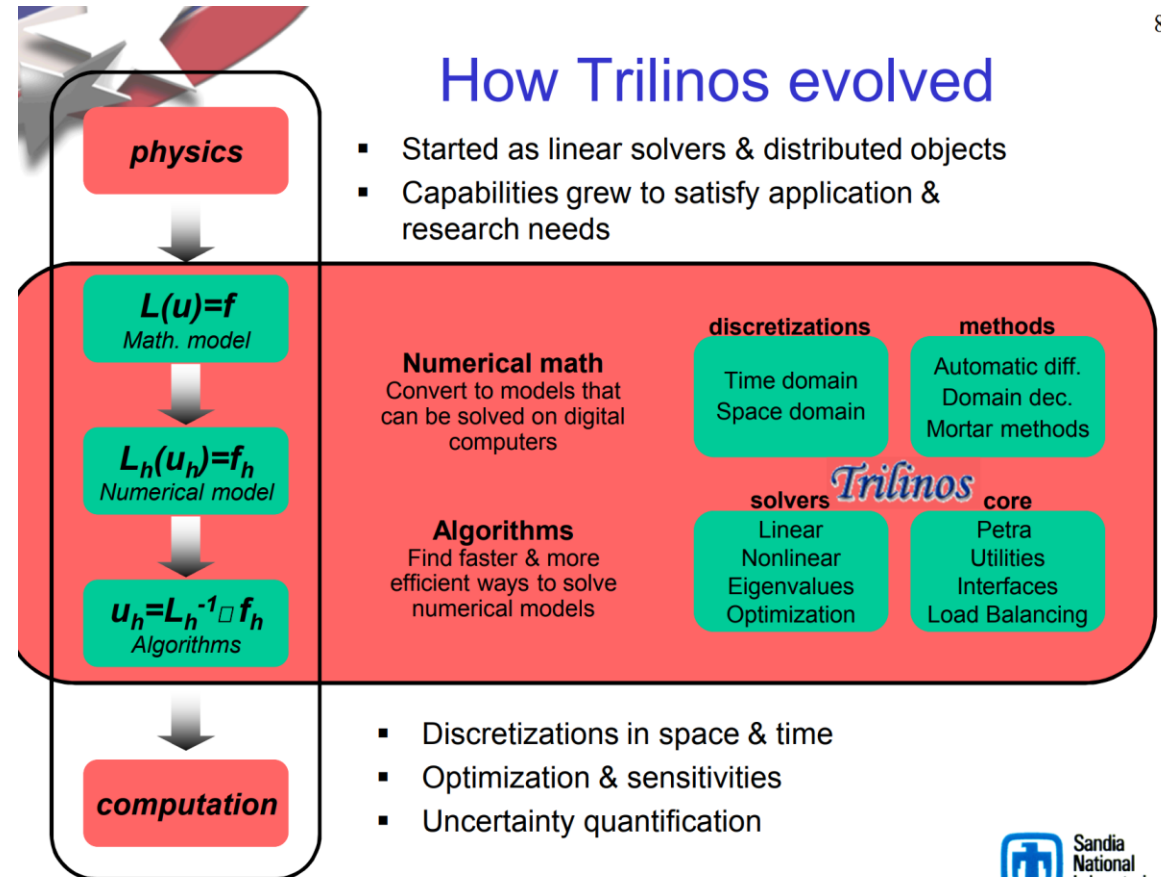
## DoD is lagging behind!





# Sandia's Trilinos Framework

• One of the main challenges for any software development is keeping the computer code up-to-date with the advancement in applied mathematics, software and hardware development in computational science and engineering. Realizing the challenge, the Computer Science Research Institute (CSRI) group at Sandia National Laboratories (Sandia) has developed and continues to develop scalable solver algorithms and software through next-gen (exa-scale, peta-scale, extreme-scale, etc.) computing investment. The project is called Trilinos project.

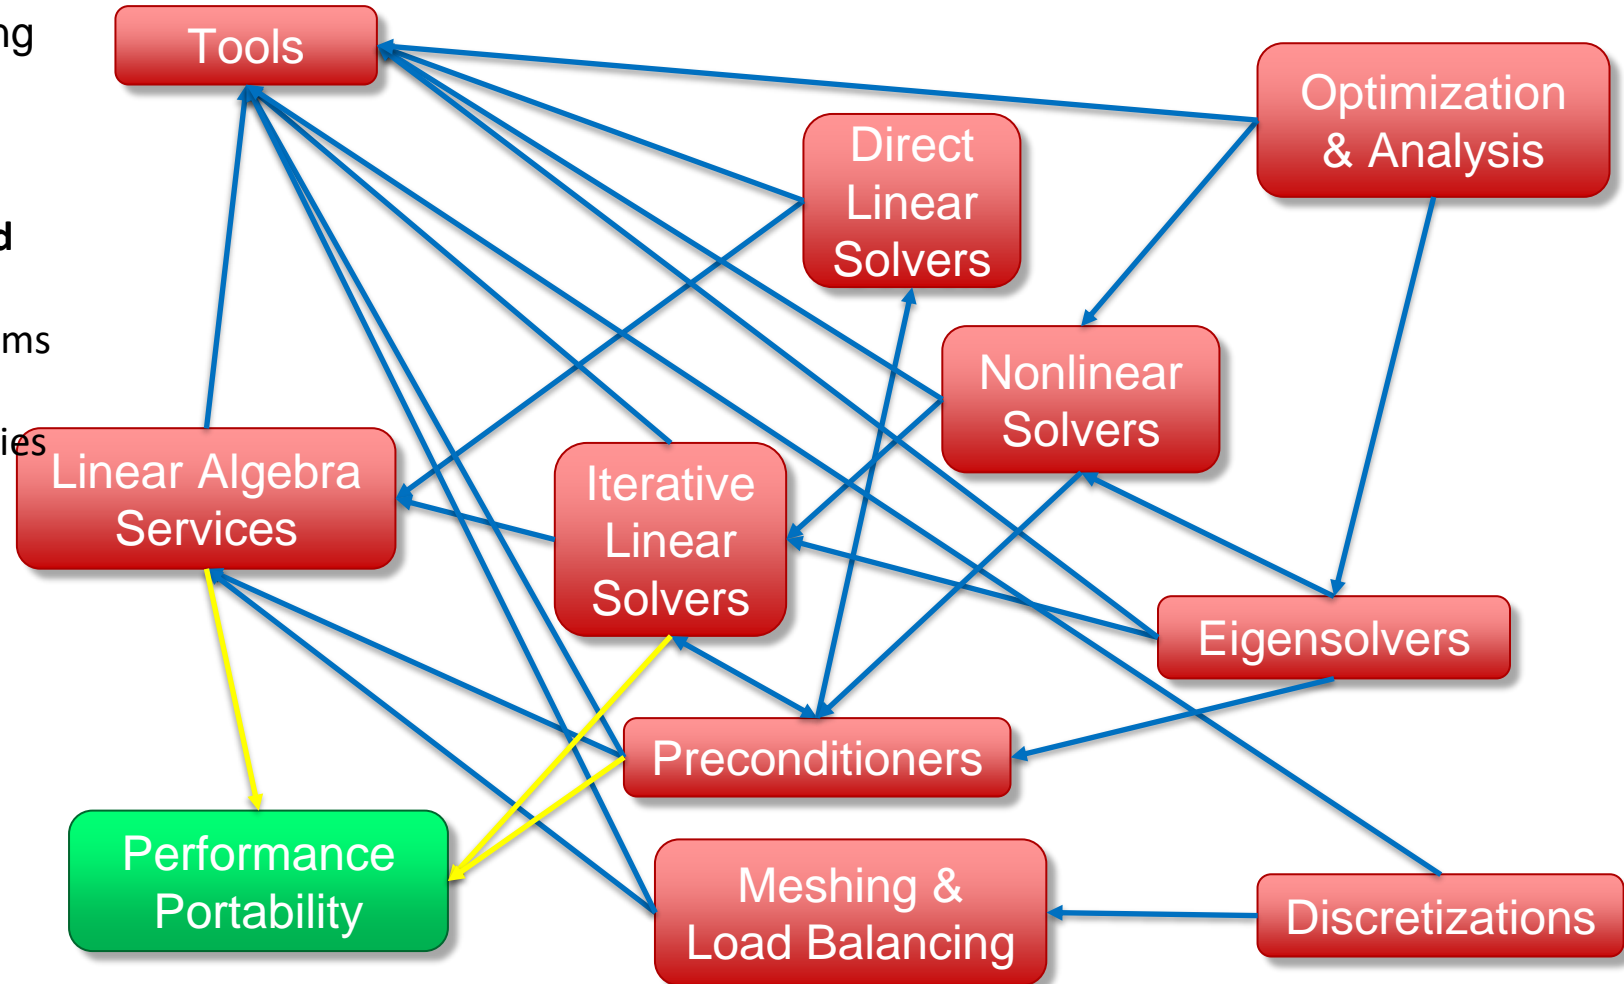


# Our Exascale Computing Efforts with Trilinos (Sandia)

- Object-oriented software framework for...
- Solving big complex science & engineering problems
- More like LEGO™ bricks than Matlab™

## Provides the state-of-the-art preconditions and linear solver libraries

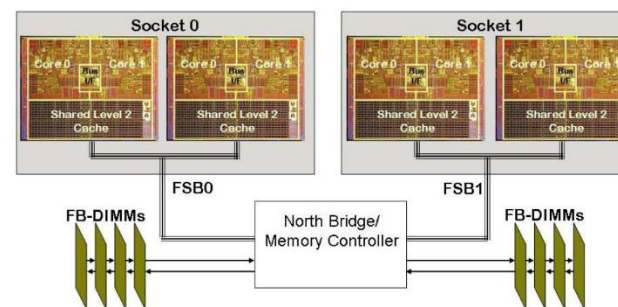
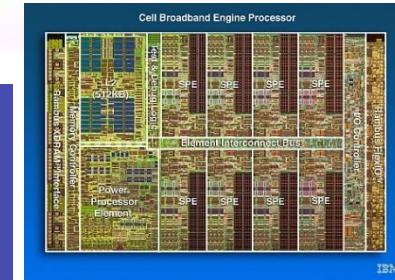

- demonstrate scalability on current HPC systems
- illustrate plans for continued maintenance
- include support for new hardware technologies



# Target Platforms



- Desktop: Development and more...
- Capability machines:
  - Redstorm (XT3), Clusters
  - Roadrunner (Cell-based).
  - Multicore nodes.
- Parallel software environments:
  - MPI
  - UPC, CAF, threads, vectors,...
  - Combinations of the above.
- User “skins”:
  - C++/C, Python
  - Fortran.
  - Web, CCA.



# Trilinos Package Advancement

Category	1 <sup>st</sup> Generation	2 <sup>nd</sup> Generation	3 <sup>rd</sup> Gen
Linear Algebra Services	<b>Epetra</b> , EpetraExt, Komplex	<b>Tpetra</b> , Xpetra, Domi, RTOp, Thyra	<b>Tpetra</b> , Xpetra
Tools	Teuchos, Triutils, Galeri, Optika, Trios	Teuchos, Sacado, Trios	
Direct Linear Solvers	<b>Amesos</b> , Pliris	<b>Amesos2</b>	
Iterative Linear Solvers	<b>AztecOO</b>	<b>Belos</b> , Stratimikos	<b>Belos</b>
Preconditioners	<b>IFPACK</b> , <b>ML</b>	<b>IFPACK2</b> , <b>MueLu</b> , ShyLU	<b>IFPACK2</b> , MueLu
Nonlinear Solvers	NOX, LOCA	NOX, LOCA	
Eigensolvers	Anasazi	Anasazi	
Optimization & Analysis	MOOCHO	MOOCHO, OptiPack, Phalanx, Piro, ROL	
Meshing & Load Balance	STK, Zoltan, Isorropia, Mesquite, Moertel	STK, Zoltan2, Pamgen	
Discretizations	Intrepid, Shards, Rythmos	Intrepid, Shards, Tempus	
Performance Portability			<b>Kokkos</b>



# MFiX-Trilinos: Advanced Linear Solver

MFiX (developed by NETL)

- Model multiphase physics
- Widely used by the fossil fuel reactor communities and beyond
- can significantly reduce time & cost to design a reactor

However

- **Computational expense** for most **practical applications** can make it **impractical**
- Limited software capabilities
  - Linear solver, MPI-X, UQ, etc.
- Can result in poor convergence especially in complex non-linear problems

But, could be made more practical if we could significantly reduce time-to-solution by

- Effectively exploiting HPC systems (massively parallel computers, GPUs, multithreading..)
- **Leveraging** state-of-the-art **preconditions and linear solver libraries**
- Providing a long-term portable and scalable software development framework

$$\mathbf{Ax} = \mathbf{b}$$

$$\epsilon_g + \sum_{m=1}^M \epsilon_{sm} = 1$$

$$\frac{\partial}{\partial t}(\epsilon_g \rho_g) + \nabla \cdot (\epsilon_g \rho_g \vec{U}_g) = R_g$$

$$\frac{\partial}{\partial t}(\epsilon_g \rho_g \vec{U}_g) + \nabla \cdot (\epsilon_g \rho_g \vec{U}_g \vec{U}_g) = -\epsilon_g \nabla P_g + \nabla \cdot \tau_g - \sum_{m=1}^M I_{gsm} + \epsilon_g \rho_g \vec{g}$$

$$\epsilon_g \rho_g C_{pg} \left( \frac{\partial T_g}{\partial t} + \vec{U}_g \cdot \nabla T_g \right) = \nabla \cdot \vec{q}_g - \sum_{m=1}^M H_{gsm} - \Delta H_{rg} + H_{wall}(T_{wall} - T_g)$$



# Goal and Objectives

---

The technical goal of this project is to develop, validate and implement **advanced linear solvers** to replace **MFiX's existing linear solvers**. This goal will be achieved by integrating **Trilinos**. The project will **demonstrate scalability** of the Trilinos- MFiX interface on various high-performance computing (HPC) facilities including the ones funded by the Department of Energy (DOE).


The expected results of the project will be **reduction of computational time** when solving complex gas-solid flow and reaction problems in MFiX, and reduction in time and cost of adding new algorithms and physics based models into MFiX

## Objectives

- Create a framework to integrate MFiX with Trilinos linear solver packages
- Validate MFiX suites of problems on HPC systems with and without GPU acceleration
- Evaluate the performance



A language independent interface to  
integrate legacy codes



MFiX wrapper

Interpret matrix  
structure, Implement  
CRS scheme

Fortran wrapper

Semantic for  
memory  
references

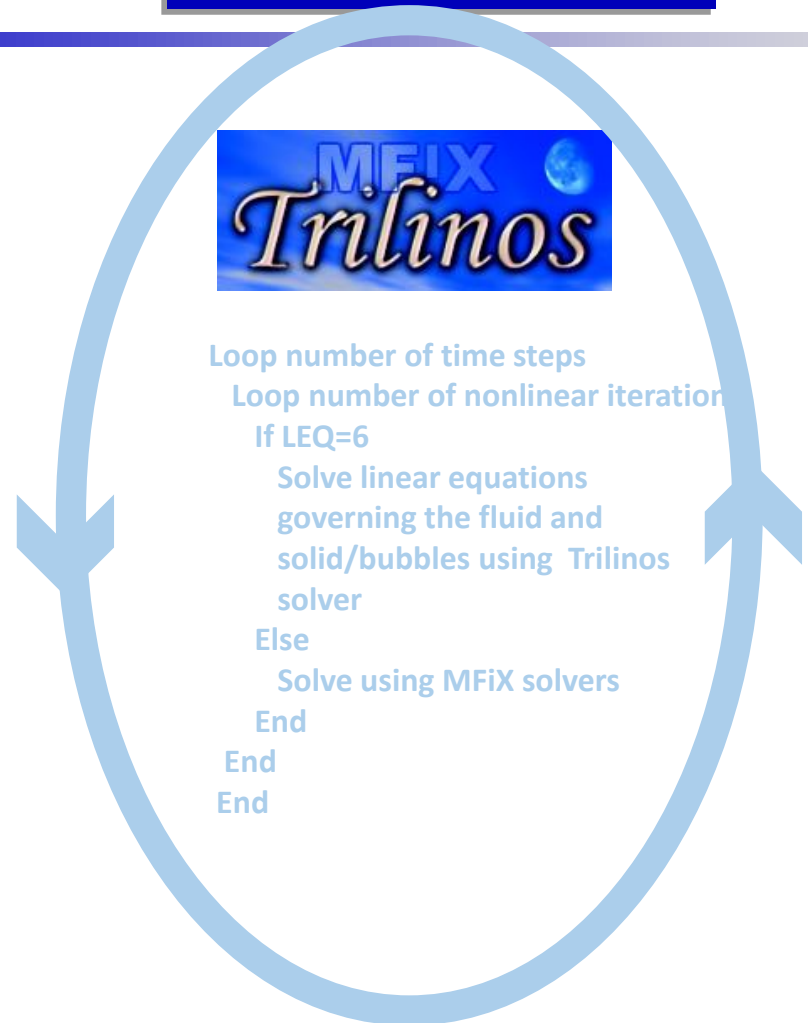

Transfer x

C wrapper

Communicator,  
Interpret  
Polymorphism  
representations

C++ wrapper

Create  
Epetra\_MAP,  
Fill A & b



# C++ Wrapper

---

```
Tpetra_Map map(numGlobalElements, numMyElements, indexBase, comm)

Tpetra_crsMatrix A(map,7);

Tpetra_multivector x (map,1);
Tpetra_multivector b (map,1);

for (LO i = 0;i < static_cast<LO> (numMyElements); ++i) {
    Values = Anew[][];
    Indices = pos[][];
    A->insertGlobalValues (gblRow, NumEntries, Values, Indices);
}
A->fillComplete (map,map);
for (LO i= 0; i < static_cast<LO> (numMyElements); ++i) {
    const GO gblRow = map->getGlobalElement (i);
    b->sumIntoGlobalValue(gblRow, 0, Bn[i]);
}

Tpetra_LinearProblem problem(&A, &x, &b);


Problem->setRightPrec (plistp);

belos_bicgstab_manager_type solver(Problem, plists));

solver->solve();
```



# Flow in a fluidized bed



# Performance - MFiX-Trilinos vs MFiX

## 3D Bubbling flow Problem

Case 1: Mesh Size = 10M

Case 2: Mesh size = 200M

## Computers:

**Stampede:** Texas Advanced Computing Center (TACC)

**Bridges:** Pittsburgh Supercomputing Center (PSC)

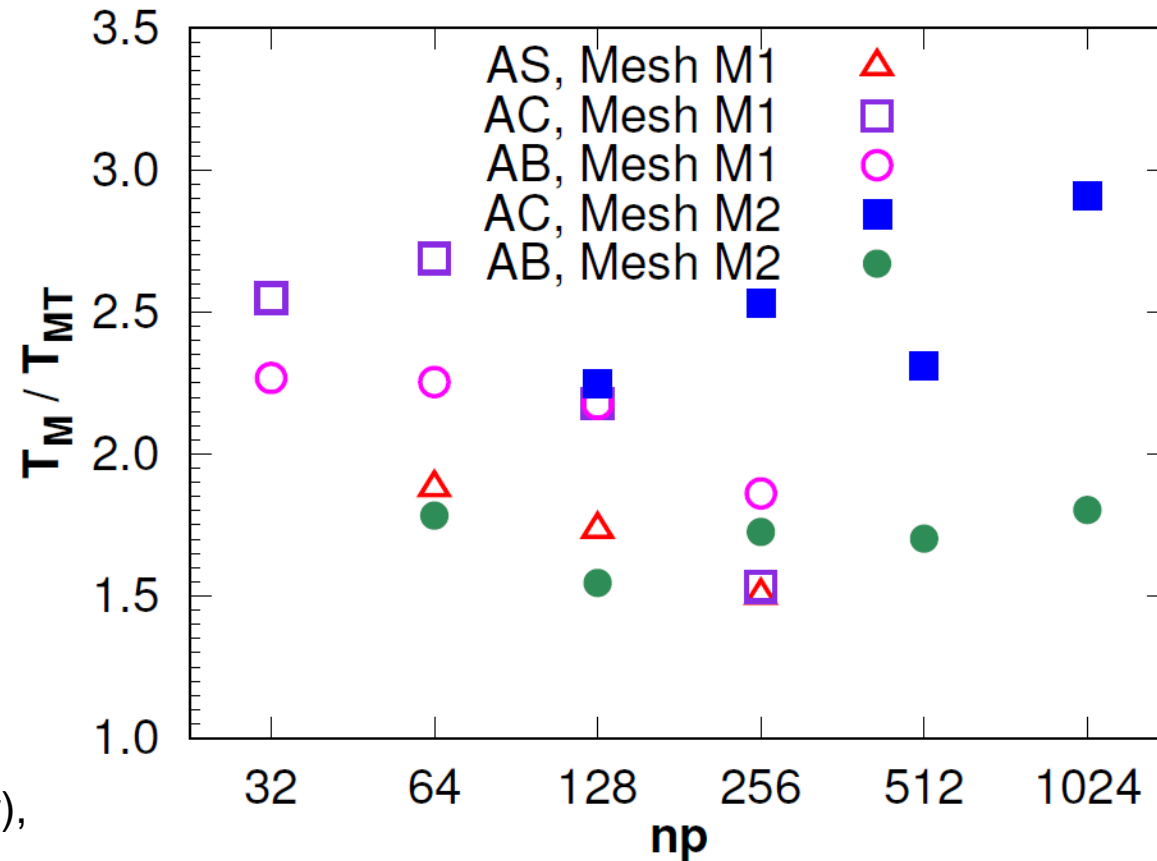
**Comet:** San Diego Supercomputer Center (SDSC)

Various computer architectures used for the performance analysis study

	Stampede (AS)	Comet (AC)	Bridges (AB)
Model (Intel Xeon)	E5-2680 2.7GHz	E5-2695 2.5GHz	E5-2695 2.30 GHz
Cores per socket	8	12	14
Sockets	2	2	2
L1 cache (KB)	32	32	32
L2 cache (KB)	256	256	256
L3 cache (KB)	20480	30720	35840
RAM (GB)	32	128	130



# Various supercomputers



## Trilinos

Solvers: BiCGStab  
Packages: Tpetra  
(obj), Belos (solver),  
MueLu (pre)

PC: Jacobi



# First and second generation solver stacks

## Second generation

Iterative Solvers: BiCGStab2, GMRES2

Packages: Tpetra (obj), Belos (solver), MueLu

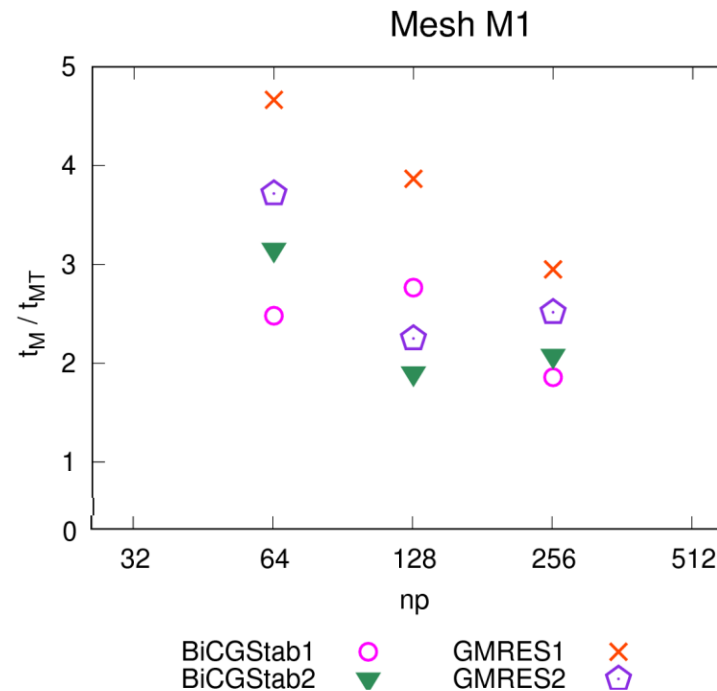
PC: Smoothed Aggregation

## First Generation

Iterative Solvers: BiCGStab1, GMRES1

Packages: Epetra (obj), Aztec(solver), ML

PC: Smoothed Aggregation



# Research Focus (with George Karniadakis, Brown)

## Physics (PDEs) & data-driven (AI/ML) Analysis for Complex Systems

Experimental Data (Physical and/or Numerical)

Lots of experiments

Data/Experiments

Physics/M&S

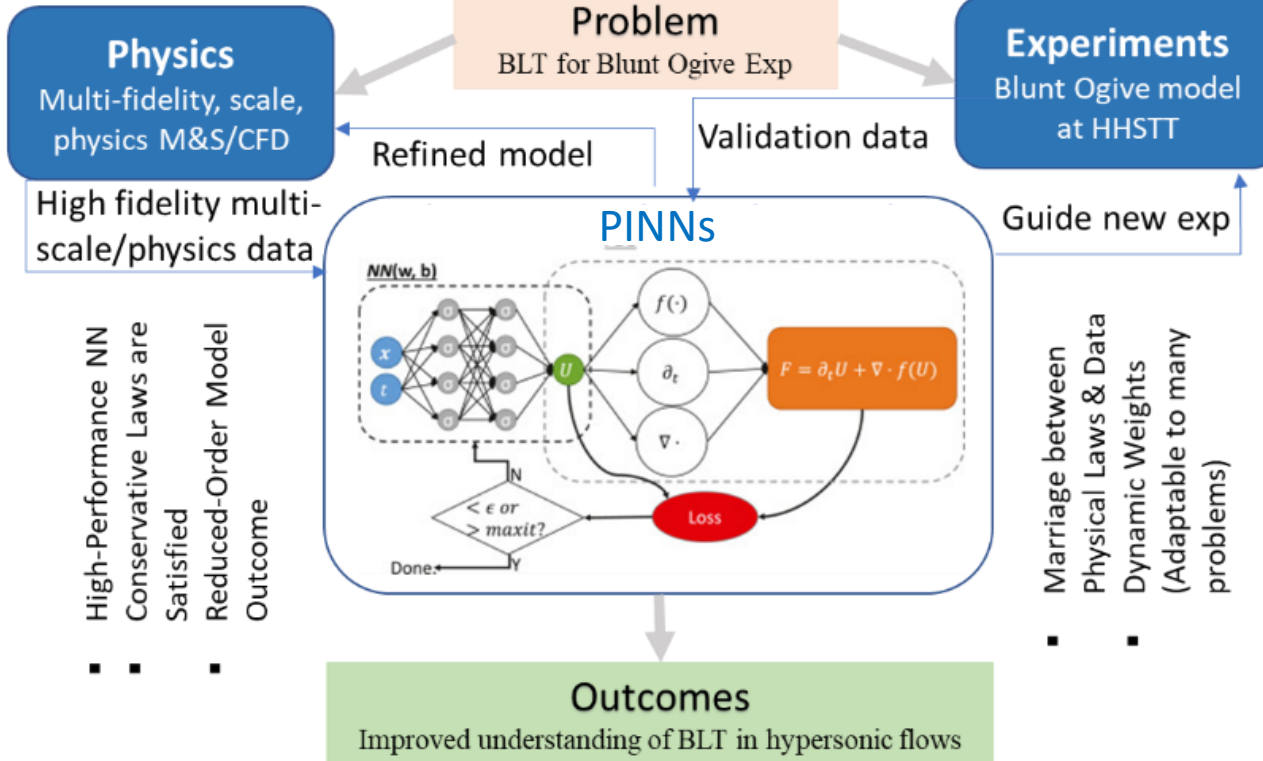

Big data

Proposed approach

Physics-Informed Neural Networks (PINNs)


### Pros & Cons

- + Less Expensive
- + Can emulate any environmental conditions
- + In-depth analytics
- Limited by fidelity & computational resources
- Computational & algorithmic limitations



- High-Performance NN
- Conservative Laws are Satisfied
- Reduced-Order Model Outcome

- Marriage between Physical Laws & Data
- Dynamic Weights (Adaptable to many problems)



### Pros & Cons

- + Real life representation
- + What you see is what you get
- Expensive
- Limited fidelity or sometimes impossible
- Operation or environmental hazard
- Measurement Limitations

Small or no data

Data/Experiments

Physics/M&S

Lots of Physics

Partial Differential Equations



# High Performance Computing (HPC)

## Advanced Modeling & Simulations, Exascale Computing

- Essential for a thriving, innovation-based economy
  - Huge opportunities at affordable price
  - Expand our innovation pipeline
  - Companies like Google, Amazon, Facebook, GE has be mostly reliant on imported talent for such needs
- Sustained government investment required for advancing U.S. leadership
  - Scaling back => putting economic power, ability to innovate & national security at risk
  - Failing to adapt disruptive innovation => leadership can slip away
- A force multiplier for innovation and efficiency
  - Good return on the investment (\$1 on HPC => \$515 overall or \$1,205 for Govt)
  - A lot of profit/cost savings (\$1 on HPC => \$52 overall or \$141 for Govt)
  - Free to us as DoD HPC users (Seed account: 250,000hrs x \$0.12/core-hr<sup>#</sup> = \$30,000 +software cost)
- Are we ready? – NOT Really!
  - Education & research resources
    - Scarce both internally at the University & externally (e.g. very difficult to research grants from funding agencies)
    - Inherently interdisciplinary but historically organization structures are designed to discourage such researchers
  - Trained human resources
    - engineering students get scared of PDEs, Numerical Methods, Programming (even at scripting level -- forget about HPC!), ...
    - Insufficient number of courses and organized education materials (such 3 credit hour courses)
  - Need for a lot advocacy and perhaps seedling support at university if we're serious
    - Decision makers (internal – deans, VPs, etc /external - PM/PD, Lab admins, etc) --- China says “...will demonstrate the world’s first exascale supercomputer

#### References:

\* <https://science.energy.gov/~media/ascr> (IDC report)

# <https://fcw.com/articles/2017/05/03/comment-hpc.aspx>

# <https://blog.rescale.com/the-real-cost-of-high-performance-computing/>  
(Equipment + Electricity + Labor + Facilities cost)

# Acquisition of an Advanced Computing Testbed for HPC, AI/ML, M&S & Digital Twin Research and Education

Amount: ~\$1M, 2024-25 (PI: Kumar, Co-PI: Yang)

Funding Agency: DOD

Award # W911NF2410267

## Infrastructure

- HPC AI/ML Testbed
- High end visualization
- Digital twin equipment




## Curricular:

- Multiphysics Computing
- Scalable & Portable Computing
- AI & Data-driven engineering




## Research:

- Digital-twin
- Modeling and Simulations
- Advanced Computing



# SciML



# Concluding Remarks

---

- **HPC Impact:** Advanced linear solvers in Trilinos significantly enhance computational capabilities in scientific research.
- **Improved Performance:** Integration with MFiX reduces computational time and costs, making simulations more feasible.
- **Future-Proofing:** Trilinos adapts to emerging technologies like GPUs, ensuring long-term relevance.
- **Collaborative Success:** Interdisciplinary teamwork is crucial for developing scalable, efficient software solutions.
- **Ongoing Challenges:** Continued R&D is essential to further improve software scalability and performance.

