CM³C

Center for Multiscale Modeling of Multiphase Combustion

AST Meeting

November 3-4, 2026

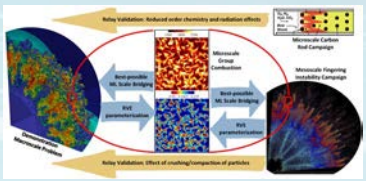

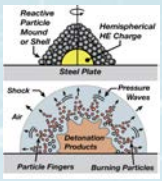

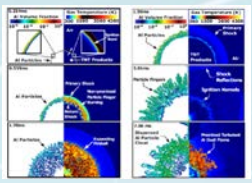

MFC + Group Combustion

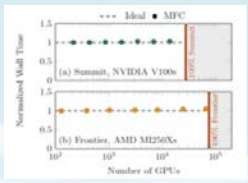

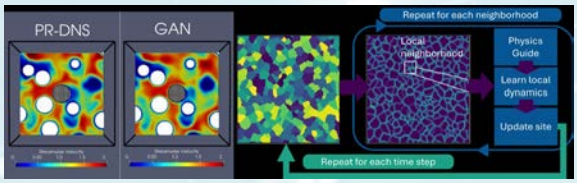

GPU + ML on Emerging Architecture

National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

Center for Multiscale Modeling of Multiphase Combustion (CM^{3}C)

TST Review Agenda — Day 1

Tuesday/Wednesday, November 3rd & 4th, 2026

Time	Topic	Presenter	Room
8:00–8:50 AM	Breakfast - All Participants		Dogwood Lane
	Private Breakfast for Review Team		Cedar
9:00–9:10 AM	Introductions		Dogwood
9:10–10:10 AM	T1. Center Overview	Balachandar	Dogwood
10:10–10:20 AM	Discussion		Dogwood
10:20–10:30 AM	Break		Dogwood Lane
10:30–11:00 AM	T2. Annual Prediction Status — Integration/workflow of required physics modules/codes and scaling	Houim	Dogwood
11:00–11:45 AM	T3. Integration	Bryngelson	Dogwood
11:45–12:00 PM	Discussion		Dogwood
12:00–12:45 PM	Lunch - All Participants		Dogwood Lane
	Private Lunch and Review Team Caucus		Cedar
12:50–1:50 PM	T4. Computer and Data Science, including AI/ML	Harley / Zare	Dogwood
1:50–2:00 PM	Discussion		Dogwood
2:00–3:00 PM	T5. Experiments and V&V-UQ	Guildenbecher / Travis / Kim	Dogwood
3:00–3:10 PM	Discussion		Dogwood
3:10–3:15 PM	Short Break for Students to Line up		Dogwood Lane
3:15–5:15 PM	Poster Session	Students	Dogwood Lane
5:15–6:30 PM	Review Team Caucus and Generation of Questions for the Center		Cedar
6:30–8:30 PM	Social & Dinner with All Participants		PDR[1]

[1] Private Dining Room

Center for Multiscale Modeling of Multiphase Combustion (CM^{3}C)

TST Review Agenda — Day 2

Tuesday/Wednesday, November 3rd & 4th, 2026

Time	Topic	Presenter	Room
8:00–8:45 AM	Breakfast - All Participants		PDR[2]
	Private Breakfast for Review Team		Cedar
8:45–9:45 AM	T6. Science Highlights	Students	PDR[2]
	MFC for Particle Resolved Simulations	Daniel Vickers	
	MFC for Euler-Lagrange Simulations	Thierry Daoud	
	N-Particle Flow-Prediction: Theoretical Foundations and Rigorous Potential Flow Testing	Davin Anderson	
	N-Particle Flow-Prediction: Data-Driven Self-Consistent Machine Learning	Raul Valle	
9:45–10:30 AM	Center Response to Questions Raised the Previous Day	Balachandar	PDR[2]
10:30–10:45 AM	Discussion		PDR[2]
10:45–11:00 AM	Break		PDR[2]
11:00–11:30 AM	T7. Plans for the Upcoming Year and Beyond	Jackson	PDR[2]
11:30 AM–3:30 PM	Private Lunch and Review Team Caucus		Cedar
2:00–2:30 PM	Discussions between Center Management	Balachandar / Jackson	PDR[2]
3:30–4:00 PM	Review Team Summary of Findings		PDR[2]
4:00 PM	Review Ends	—	

[2]Private Dining Room

Center for Multiscale Modeling of Multiphase Combustion (CM^3C)

AST Review November 3-4, 2026, Attendee List

Faculty

S. Balachandar "Bala"	University of Florida	bala1s@ufl.edu
Spencer Bryngelson	University of Texas at Austin	shb@utexas.edu
James Fairbanks	University of Florida	fairbanksj@ufl.edu
Daniel Guildenbecher	Purdue University	dguilden@purdue.edu
Joel Harley	University of Florida	joel.harley@ufl.edu
Ryan Houim	Purdue University	rhouim@purdue.edu
Tom Jackson	University of Florida	tlj@ufl.edu
Nam-Ho Kim	University of Florida	nkim@ufl.edu
Alina Zare	University of Florida	azare@eng.ufl.edu
ST Thakur	University of Florida	sst@ufl.edu
Claire Grégoire	University of Florida	c.gregoire@ufl.edu

Review Team

John Hewson (Chair)	SNL	jchewso@sandia.gov
Bradford Durant	LANL	bdurant@lanl.gov
Ben Isaac	LLNL	isaac4@llnl.gov
Fady Najjar	LLNL	najjar2@llnl.gov
David Noble	SNL	drnoble@sandia.gov
Kathy Prestridge	LANL	kpp@lanl.gov

Alliance Strategy Team (AST)

Tim Germann (Chair)	LANL	tcg@lanl.gov
Coleman Alleman	SNL	callema@sandia.gov
Joshua Dick	NNSA/Graduate Fellow	joshua.dick@nnsa.doe.gov
Erik Draeger	LLNL	draeger1@llnl.gov
Sriram Swaminarayan	LANL	sriram@lanl.gov

Tri-Lab Sponsor Team (TST)

Flint Pierce (Chair)	SNL	fpierce@sandia.gov
Georges Akiki	LLNL	akiki1@llnl.gov
Angela Herring	LANL	angelah@lanl.gov
Christopher Malone	LANL	cmalone@lanl.gov
Daniel Richardson	SNL	drrich@sandia.gov
Carol Woodward	LLNL	woodward6@llnl.gov

Center for Multiscale Modeling of Multiphase Combustion (CM^3C)

Research Staff

BC Kim	University of Florida	kim.byeongch@ufl.edu
Erica Lindbeck	University of Florida	elindbeck@ufl.edu
Colton Willhardt	Purdue University	cwillhar@purdue.edu
Jacob Posey	Purdue University	jwposey@purdue.edu

Administration Staff

Linda Curwick	University of Florida	lindacurwick@ufl.edu

Students

Aayushma Aashu	University of Florida	aayushma.aashu@ufl.edu
Davin Anderson	University of Florida	andersondavin@ufl.edu
Avery Couture	University of Florida	avery.couture@ufl.edu
Thierry Daoud	University of Florida	t.daoud@ufl.edu
Juan Diaz	University of Florida	jdiaz1@ufl.edu
Joshua Gillis	University of Florida	joshuagillis@ufl.edu
Christian King	University of Florida	christianking@ufl.edu
Cameron Lott	University of Florida	lottc@ufl.edu
Alex Millauer-Hsu	Purdue University	amillaue@purdue.edu
Fahad Nabid	University of Florida	fahadnabid@ufl.edu
Cohen Nunes	Purdue University	cnunes@purdue.edu
Gregory O'Brien	Purdue University	greg99@purdue.edu
George Rauta	University of Florida	grauta@ufl.edu
Julian Reyes	University of Florida	reyes.julian@ufl.edu
Isabella Rivera-Sigmann	University of Florida	il.rivera@ufl.edu
Chrissie Spangler	Purdue University	spanglc@purdue.edu
Jacob Stewart	University of Florida	jr.stewart@ufl.edu
Isaly Tappan	University of Florida	i.tappan@ufl.edu
Andreas Thompson	University of Florida	thompsona1@ufl.edu
Rebekah Travis	Purdue University	travisr@purdue.edu
Mike VandenBoom	University of Florida	m.vandenboom@ufl.edu
Raul Valle	University of Florida	rvalle1@ufl.edu
Daniel Vickers	Georgia Tech	dvickers6@gatech.edu
Benjamin Wilfong	Georgia Tech	bwilfong3@gatech.edu

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

T1: Center Overview

Principal Investigator: S. Balachandar

Department of Mechanical and Aerospace Engineering

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

UF UNIVERSITY of FLORIDA

Center Related Questions

1. Review the **goals and objectives** of your Center, including your overarching application
2. Provide a one-viewgraph **simulation roadmap** that indicates where you are at present and how you will reach your final predictive simulation in year five
   - This roadmap should clearly show how all of the components of your program come together to advance your overall prediction goal, and be referred to in subsequent talks and posters.
3. Identify a few (3-5) **key highlights** from the past year – what are you most excited about? (With details to be given in a subsequent talk/poster.)
4. Describe how ASC TriLab (and other) **supercomputing resources** were utilized, e.g., the job mix between smaller-scale single-physics, code development and testing, and fullscale predictive simulation run(s), and any issues encountered.
5. Include an **organization chart**, highlighting any changes since the proposal/site visit and staffing (including student and postdoc recruiting).
6. Discuss **student Lab internship** results and plans, and other notable Lab collaborations and interactions.
7. Describe the status of **educational programs** in computational and predictive science, scientific machine learning, applied math, and/or computer science.

Agenda to Follow

- T2 – Annual Prediction Status (Ryan Houim)
- T3 – Integration (Spencer Bryngelson)

- T4 – Data Science AI/ML (Joel Harley, Alina Zare)
- T5 – Experiments, V&V, UQ (Daniel Guildenbecher, NamHo Kim)

- Poster session (Students)
- Review Team Caucus

- T6 –Science Highlights (Daniel Vickers, Thierry Daoud, Davin Anderson, Raul Valle)
- T7 – Answers to Questions (Balachandar)
- T8 – Plans for upcoming year (Jackson)

- Discussion with Center Management
- Outbriefing

Outline

- Go over the challenges of our demonstration problem
- Revisit technical objectives and scientific goals
- Our closure modeling philosophy
- Multiscale modeling innovation by interweaving simulations and ML
- Experiments, validation, and uncertainty
- What we have accomplished so far
- What lies ahead

Our PSAAP-IV Team

CM³C Physics and Simulations


S. Balachandar *(Florida)*


Thomas Jackson *(Florida)*


Ryan Houim *(Purdue)*


ST Thakur *(Florida)*

Exascale Computing


Spencer Bryngelson *(Texas)*

Experiments


Daniel Guildenbecher *(Purdue)*

V&V UQ and Workflow


Nam-Ho Kim *(Florida)*


James Fairbanks *(Florida)*

Machine Learning


Alina Zare *(Florida)*


Joel Harley *(Florida)*

Partners


Rodney Fox *(Iowa State)*


Brian Taylor *(Eglin AFB)*


Don Littrell *(Eglin AFB)*

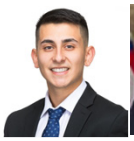
Jose Aguilar *(Eglin AFB)*


John Spencer *(NSWC Crane)*

CM³C

Our PSAAP-IV Team – Postdocs and Students


BC Kim *(Florida)*


Erica Lindbeck *(Florida)*


Jacob Posey *(Purdue)*


Colton Willhardt *(Purdue)*


Aayushma Aashu *(Florida)*


Davin Anderson *(Florida)*


Avery Couture *(Florida)*


Thierry Daoud *(Florida)*


Juan Diaz *(Florida)*


Joshua Gillis *(Florida)*


Christian King *(Florida)*


Cameron Lott *(Florida)*


Alex Millauer-Hsu *(Purdue)*


Fahad Nabid *(Florida)*


Cohen Nunes *(Purdue)*


Gregory O'Brian *(Purdue)*

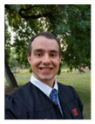
George Rauta *(Florida)*


Julian Reyes *(Florida)*


Isabella Rivera *(Florida)*


Chrissie Spanget *(Purdue)*


Jacob Stewart *(Florida)*


Isaly Tappan *(Florida)*


Andreas Thompson *(Florida)*


Rebekah Travis *(Purdue)*


Raul Valle *(Florida)*

Mike VandenBoom *(Florida)*

Daniel Vickers *(GT)*

Ben Wilfong *(GT)*

CM³C

Demonstration Problem

- Yearly integrated simulations starting from year-1
- Validation experiments & measurements at Purdue (and partners)

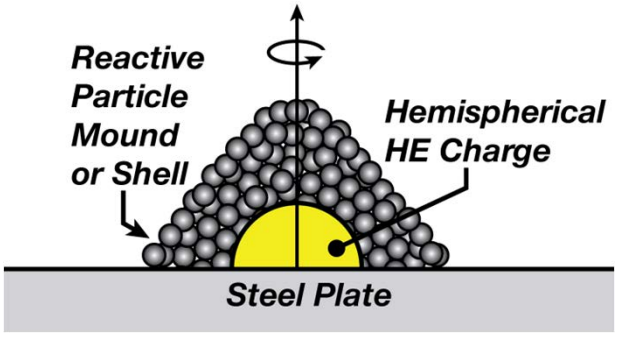

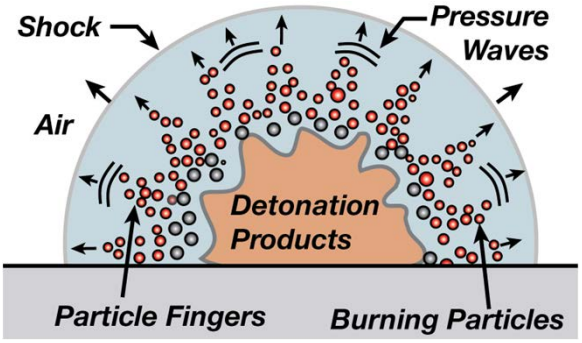

CM³C – Intellectual Objectives

Center for Multiscale Modeling of Multiphase Combustion (CM³C)

- Radically advance the understanding of *Group Combustion Physics* built upon hierarchical physics-informed ML models with uncertainty quantification
- Develop a trace-based load-balancing strategy for multiphase simulations using the proven exascale code MFC (Multicomponent Flow Code)
- Develop workflow strategies for seamless integration of experiments, simulations, ML models, and UQ across scales
- Perform validation-quality demonstration, meso and microscale experiments for relay validation and uncertainty reduction
- Educate students and postdocs in exascale simulation science and place them at NNSA laboratories

Demonstration Problem – Detonation & Dispersion

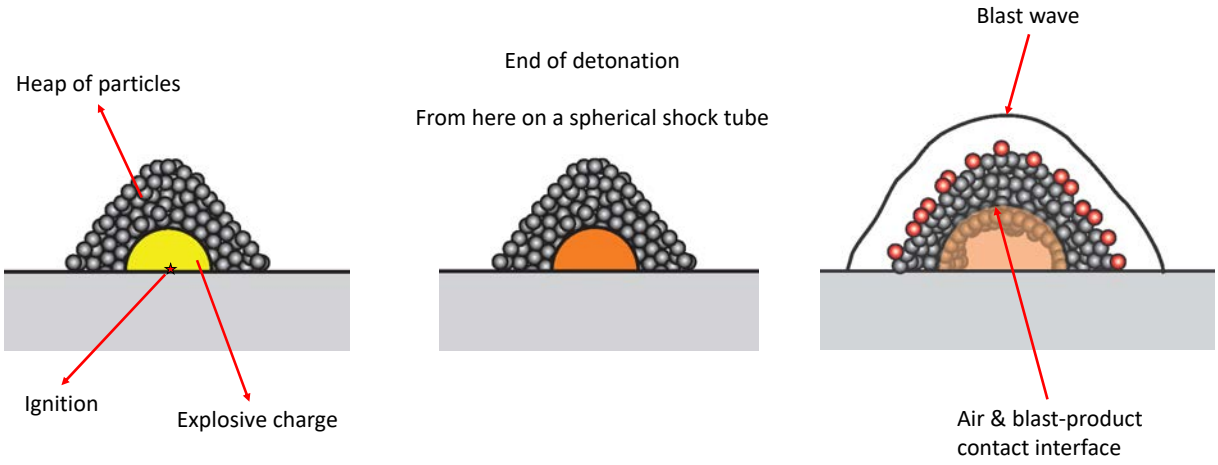

CM³C

Importance of Group Ignition & Combustion

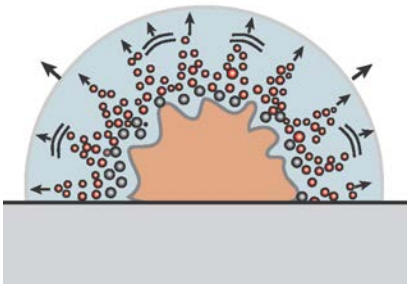

Burning particles interact at microscale

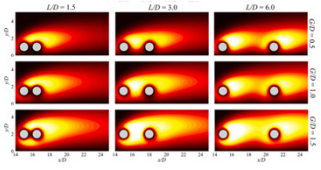

Burning particles interact at mesoscale

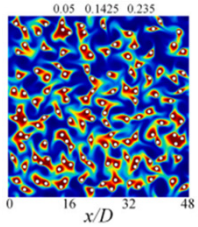

- Particles are non-uniformly distributed (fingering instability)
- It matters where the particles are (shock-heated vs det-prod)
- Ignition of one particle can trigger ignition of the cluster
- A burning particle increases temperature, but decreases oxidizer in the wake
- Burn rate of a cluster is vastly different from that of an isolated particle

CM³C

UF UNIVERSITY of FLORIDA

Global Influence – Diverse Large-scale Behaviors

Small particle-to-charge ratio

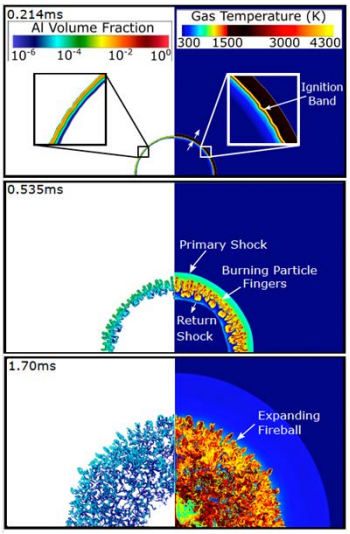

Large particle-to-charge ratio

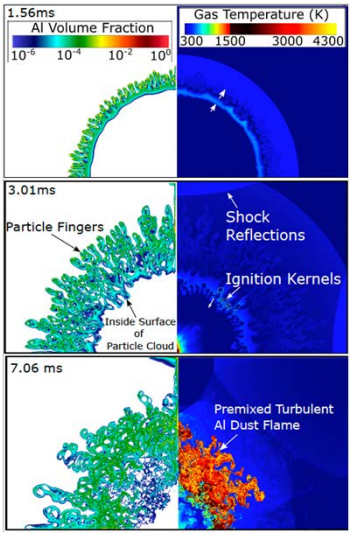

- Coarse grained Euler-Euler simulations of Al particles
- Left: small particle-to-charge ratio
  - Prompt ignition
  - Non-premixed burning
- Right: large particle-to-charge ratio
  - Delayed ignition
  - Premixed burning

CM³C

UF UNIVERSITY of FLORIDA

Technical Objectives & QoI

- Use exascale to predict (approach first-principle approach in accuracy)
  - Non-uniform distribution of explosively dispersing particles
  - Resulting spatio-temporal variation of burn-rate
  - Multiphase afterburn effect on blast dynamics
- Impact the predictive capability of engineering tools
  - With the "Best-possible" closure models

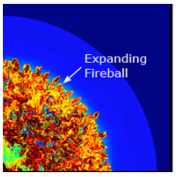

Quantities of Interest (QoI)

QoI1: Time-dependent blast location

QoI2: Time-dep pressure/impulse at probes

QoI3: Volume fraction of reactive particles

Basic metrics

QoI4: Temp and heat release rate

QoI5: Number & amplitude of instabilities

QoI6: PDFs of pressure & volume fraction

Higher-order metrics

CM³C

Proposed Strategy & Rationale

Demonstration Macroscale Problem

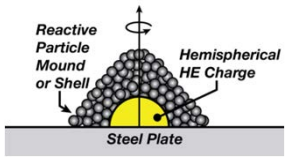

~ 3 cm Hemisphere PETN

Spherical shell of C/Mg particles
~ $O(10^{10})$ to $O(10^{11})$ particles

Initiation
Detonation
Compaction
Dispersion
Afterburn

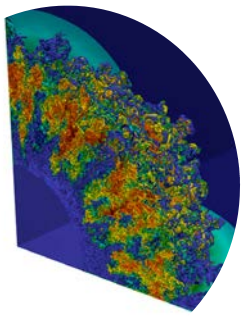

~ $(1.5\ m)^3$ volume of reactive particles

Modeling & Simulation Challenges

- Stiff chemical kinetics
  - Numerous reactions & species
  - Surface/gas-phase, premixed/non-premixed
- High-speed flow
  - Intense shocks and other discontinuities
  - Non-classical turbulence
  - Real gas EOS and properties
- Multiphase
  - Trillions of time-varying polydisperse particles
  - Direct particle-particle collisions
  - Droplet deformation and break-up

CM³C

What Can Be Simulated at Exascale

Demonstration Macroscale Problem

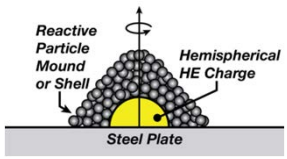

~ 3 cm Hemisphere PETN

Spherical shell of C/Mg particles
~ $O(10^{10})$ to $O(10^{11})$ particles

Initiation
Detonation
Compaction
Dispersion
Afterburn

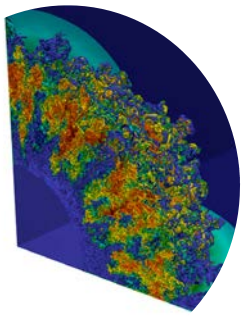

~ $(1.5\ m)^3$ volume of reactive particles

- Fully-resolved DNS
  - Full chemistry, particles, turbulence



- Euler-Lagrange (EL) LES
  - $O(10^{12})$ grid at ~300 µm resolution
  - $O(10^{10})$ to $O(10^{11})$ particles
  - O(10) to O(100) chemical reactions
  - JWL or BKW EOS, variable properties



- Euler-Euler (EE) LES/URANS
  - $O(10^{8})$ grid at 6 mm resolution
  - No particle tracking
  - …

CM³C

DNS: Direct numerical simulation; LES: Large eddy simulation; URANS: Unsteady Reynolds averaged Navier Stokes

Two Technical Objectives

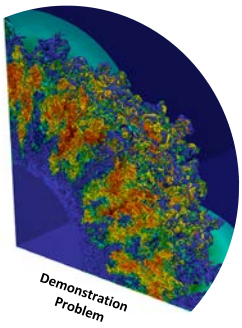

Develop "best-possible" EL closure models, so that exascale- LES is "statistically" as good as HF-PR

Develop "best possible" EE closure models to impact engineering simulations

- Fully-resolved DNS
  - Full chemistry, particles, turbulence
- Euler-Lagrange (EL) LES
  - O(10^{12}) grid at ~300 µm resolution
  - O(10^{10}) to O(10^{11}) particles
  - O(10) to O(100) chemical reactions
  - JWL or BKW EOS, variable properties
- Euler-Euler (EE) LES/URANS
  - O(10^{8}) grid at 6 mm resolution
  - No particle tracking
  - …

DNS: Direct numerical simulation; LES: Large eddy simulation; URANS: Unsteady Reynolds averaged Navier Stokes

Rigorously Derived EL Equations

Quantities computed within a RVE
Exascale EL-LES (coarse-grained)

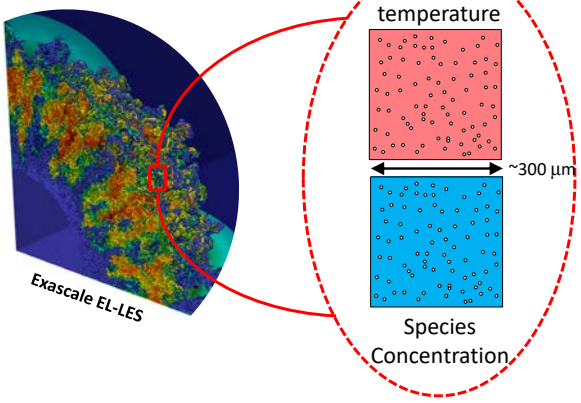

Filtered LES equations

Particle quantities (Lagrangian)

$$\bar{\mathcal{F}}\left(\bar{\mathbf{u}}(\mathbf{x},t),\mathbf{v}(j,t),F(j,t)\right)=0$$

Fluid quantities (Eulerian)

Phase-coupling closure quantities

- Force, torque on each particle
- Evaporation/burn rate of each particle

Subgrid closure quantities

- Subgrid Reynolds stress
- Subgrid energy flux
- Subgrid kinetic energy flux

RVE: Representative volume element = one (~300 µm)^3 grid cell

Closure Dependence – There is Rich Complex Information

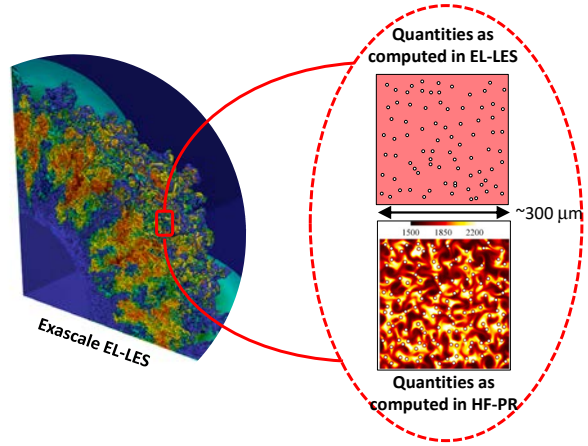

EL closure quantities modeled as functions of

$$\bar{\mathcal{F}}\left(\bar{\mathbf{u}}(\mathbf{x},t),\mathbf{v}(j,t),F(j,t)\right)=0$$

- Flow information within RVE (λ_f):
  - $\bar{\mathbf{u}}$, $\bar{T}$, $\bar{p}$, $\bar{c}_{O_2}$, …, particle volume fraction
  - Spatio-temporal gradients (i.e, $\nabla\bar{\mathbf{u}}$ and $\partial\bar{\mathbf{u}}/\partial t$)
- Particle information within RVE (λ_p):
  - Relative position, velocity, rotation
  - Temperature, burn rate, etc.,

Closure model } $(\lambda_f, \lambda_p) \rightarrow F$

CM³C

Traditional Closure Models – Greatly Simplify Dependence

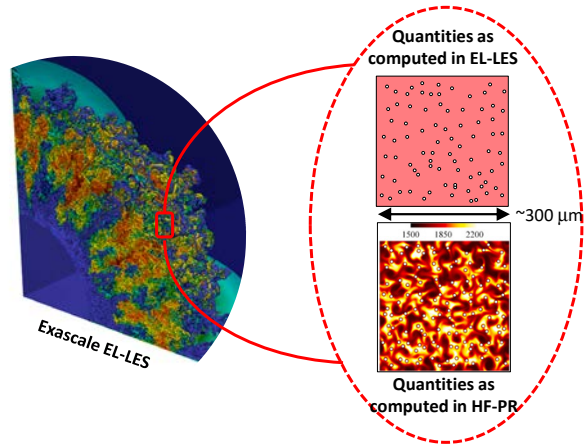

EL closure quantities

$$\bar{\mathcal{F}}(\bar{\mathbf{u}}(\mathbf{x},t),\mathbf{v}(j,t),F(j,t))=0$$

Closure model } $(\lambda_f, \lambda_p) \rightarrow F$

- Phase-coupling (drag, burn rate) models
  - Based on an isolated particle in steady uniform flow
  - Average influence of uniform distribution of neighbors
- Subgrid (Reynolds stress, flux) models
  - Assume Smagorinsky model
  - Average influence of uniform distribution of neighbors

RVE: Representative volume element = one (~300 μm)³ grid cell

CM³C

Key Principles Make Our Approach Unique

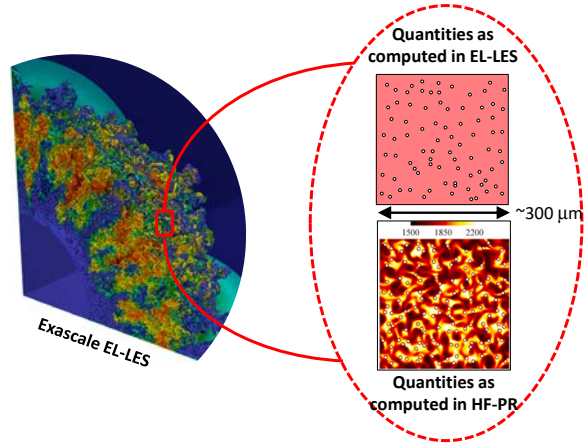

EL closure quantities

$$\bar{\mathcal{F}}\left(\bar{\mathbf{u}}(\mathbf{x},t),\mathbf{v}(j,t),F(j,t)\right)=0$$

Closure model $(\lambda_f,\lambda_p)\rightarrow F$

1) Filtered EL-LES are rigorous. So we seek

Best-possible $F(j,t)$ = Best-possible EL-LES

2) Lagrangian-based closure - $F(j,t)$

3) F has both epistemic and aleatoric uncertainty

We seek PDF of closure quantities: $P(F)$

4) We retain the full complexity of (λ_f,λ_p)

5) All closures quantities modelled together

RVE: Representative volume element = one (~300 μm)^3 grid cell

Best-Possible Closure

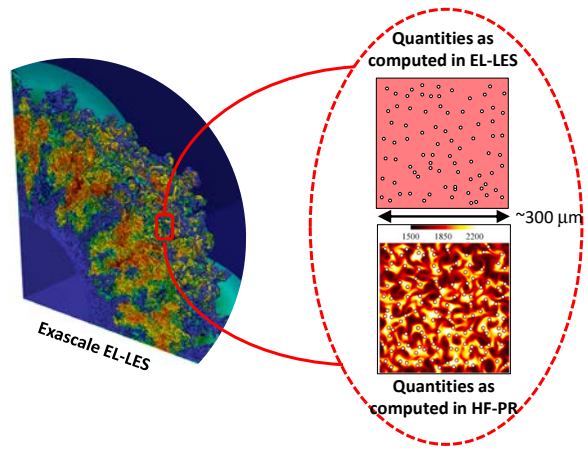

EL closure quantities

$$\bar{\mathcal{F}}\left(\bar{\mathbf{u}}(\mathbf{x},t),\mathbf{v}(j,t),F(j,t)\right)=0$$

Closure model $(\lambda_f,\lambda_p)\rightarrow P(F)$

- Best-possible closure:

  Given any (λ_f,λ_p), the predicted $P(F)$ must match the particle-resolved statistics

- Accurate subgrid flow prediction
- Must learn to predict for <u>every relevant</u> (λ_f,λ_p)

RVE: Representative volume element = one (~300 μm)^3 grid cell

Enabling Innovation – I: Flow Prediction + Closure Prediction

Flow prediction (first) ⇒ closure prediction (second)

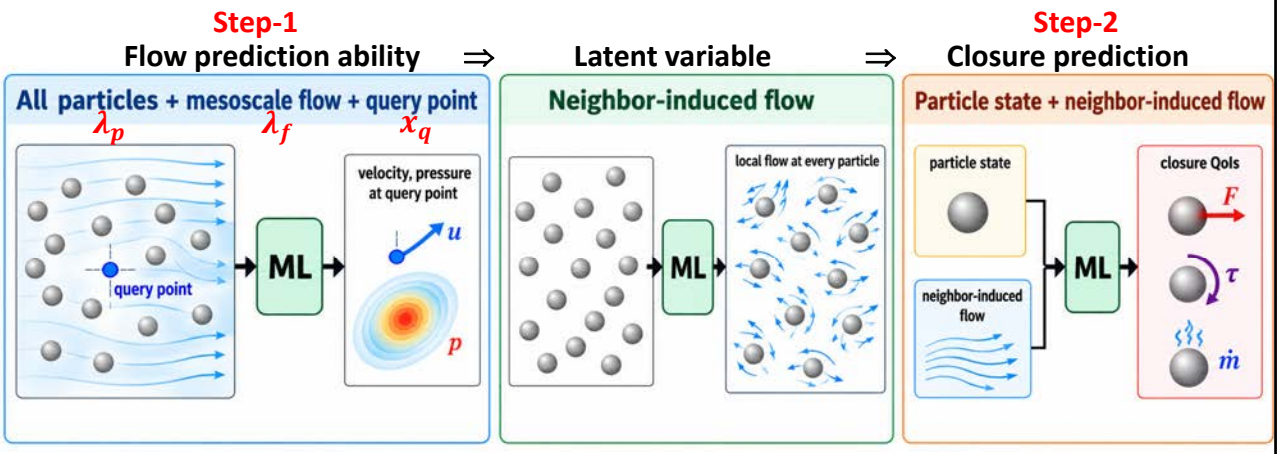

- Predicted flow yields: (i) the latent variable of all the particles, (ii) all the closure quantities

CM³C

Enabling Innovation – II: Composition AI (Physics-based)

A complex physical solution can be learned more effectively by composing simpler reusable operators than by learning one monolithic mapping

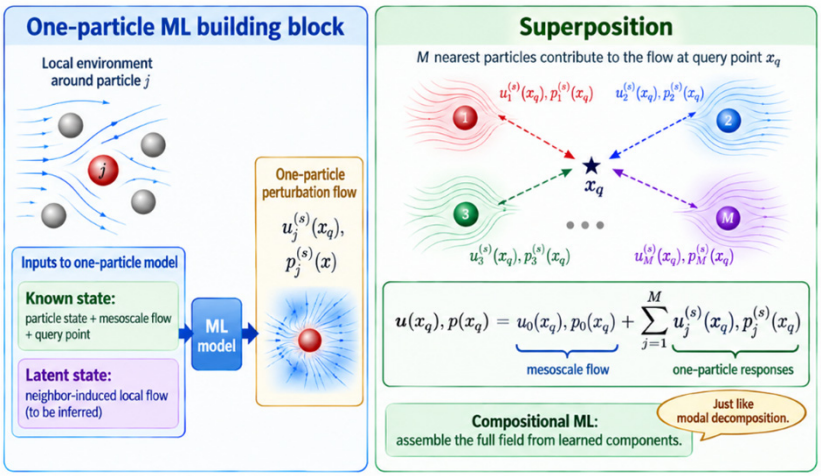

CM³C

Enabling Innovation – III: Assume Guarantee

- Neighbor-induced latent states must satisfy self-consistency
- Fixed-point iteration with embed ML training is needed

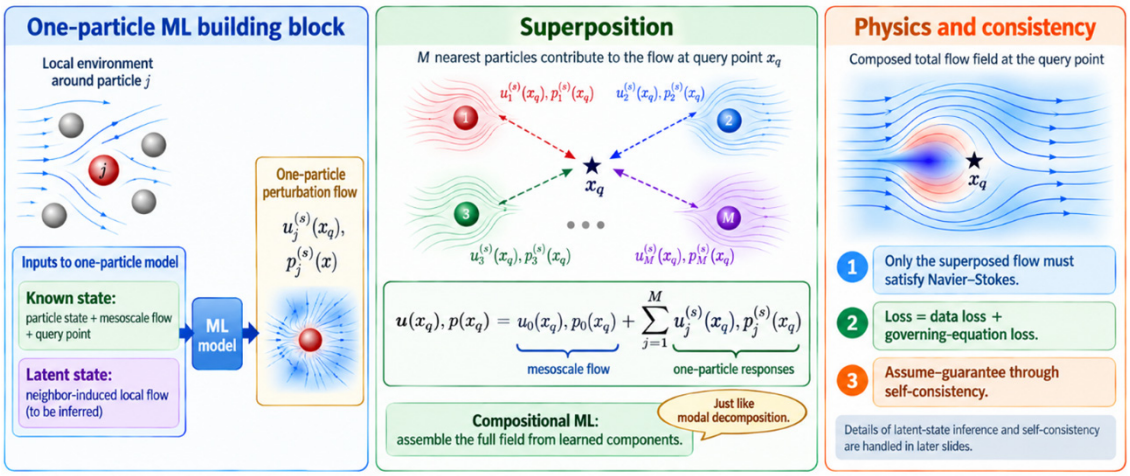

Closure Prediction Becomes a One-Particle Problem

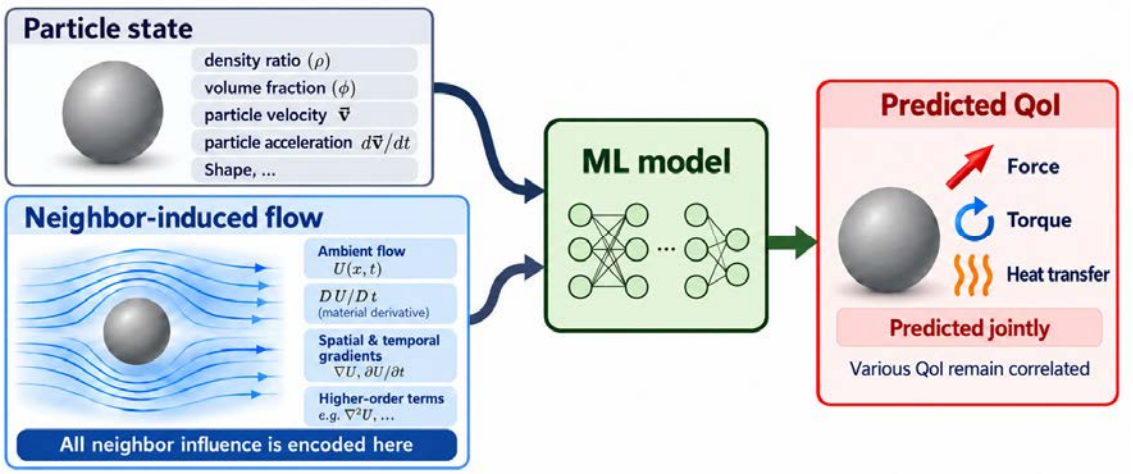

- The same predicted subgrid flow is used to consistently obtain all the closure quantities

UF UNIVERSITY of FLORIDA

How Well Superposition + Best-Possible Closure Work?

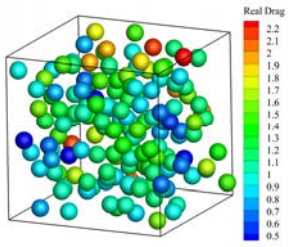

Akiki et al 2017

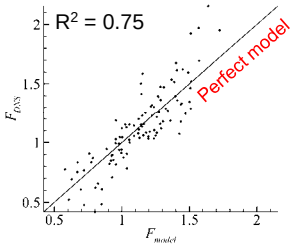

$\rho_p/\rho_f = 1.5;\ Re_p \sim 238$

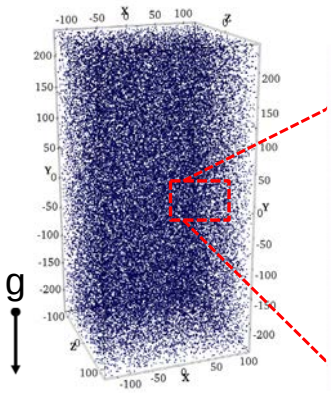

Moore & SB 2019
SB et al., 2020

❖ Geometry & Particles

- Domain: $L_x = 240, L_y = 480\ L_z = 240$

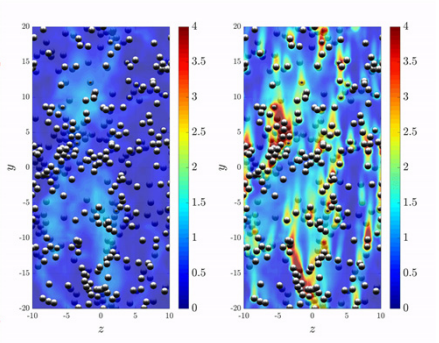

EL - uses Force prediction

With flow prediction

Effective grid $= 226$ B

CM³C

UF UNIVERSITY of FLORIDA

How Well Does This Idea Work? 3D Potential Flow Example

A cloud of evaporating droplets

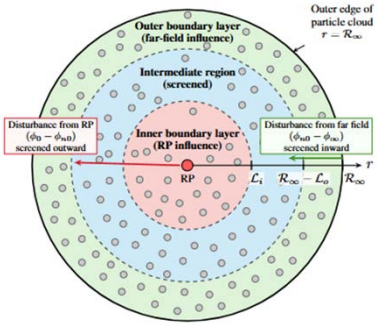

See Davin's Poster & Presentation

- One drop (RP) surrounded by a 3D cloud of random droplets at $\phi = 3\%$
- Temperature of RP and all others known
- Can we predict evaporation rate & force?
- 7500 random realizations

Can statistical closures predict each realization?

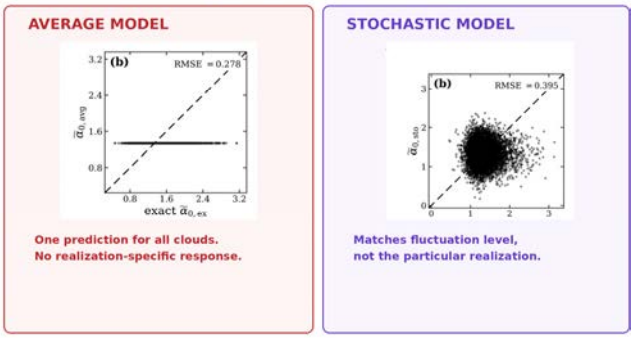

- Standard average and the stochastic models are excellent in predicting the mean/rms
- Incapable of predicting droplet-to-droplet variation

Davin, A., Jacson, T., Balachandar, S., *Flow Prediction in Hydrodynamically Coupled N-Particle Systems via Superposition of One-Particle Solutions with Potential Flow Example*, Submitted JCP, 2026

CM³C

How Well Does This Idea Work? 3D Potential Flow Example

Flow Prediction Based Deterministic Model

- One-particle Stefan flow known analytically
- Neighbor-induced flow as multipole expansion
- Each particle is influenced by K neighbors
- Coupled system solved iteratively

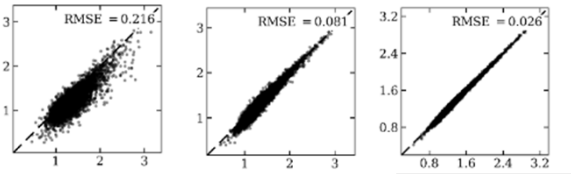

K=1 Neighbor | K=5 Neighbors | K=20 Neighbors

Predicting the neighbor-induce flow is the Key

Contours of velocity Potential

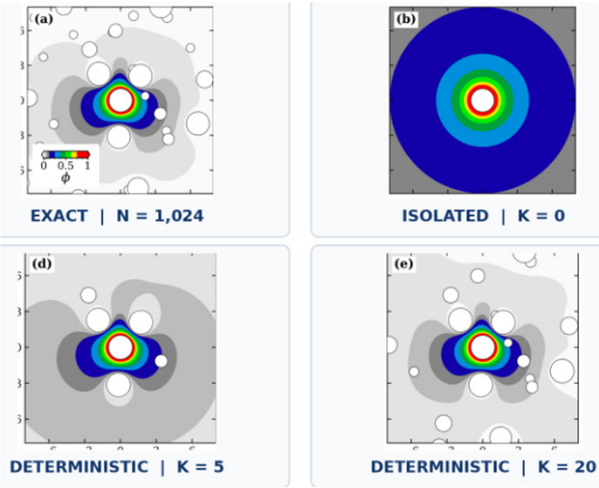

Davin, A., Jacson, T., Balachandar, S., *Flow Prediction in Hydrodynamically Coupled N-Particle Systems via Superposition of One-Particle Solutions with Potential Flow Example*, Submitted JCP, 2026

How Well Does This Idea Work? Finite Re 2D Example

Flow Prediction Based Deterministic Model

- One-particle solution (primitive) using PINN
- Neighbor-induced flow approximated as uniform flow
- Each particle is influenced by K neighbors (interaction)
- Coupled system solved iteratively

See Joel's & Raul's Presentations

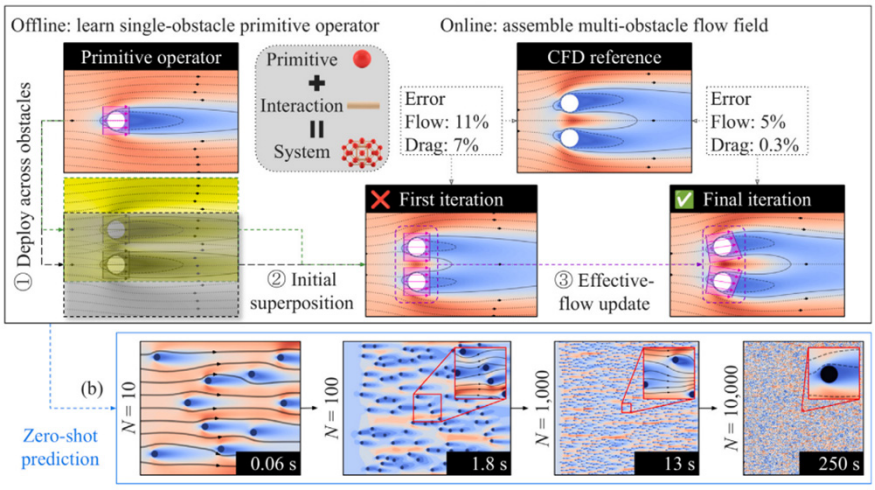

90% prediction accuracy at orders of magnitude speedup

Liu, K., Li, H., Balachandar, S., *JCP*, 2025 and Liu, K., Li, H., Balachandar, S., *Computers & Fluids*, 2025

PR Simulation ⇒ ML ⇒ EL Simulation Pipeline

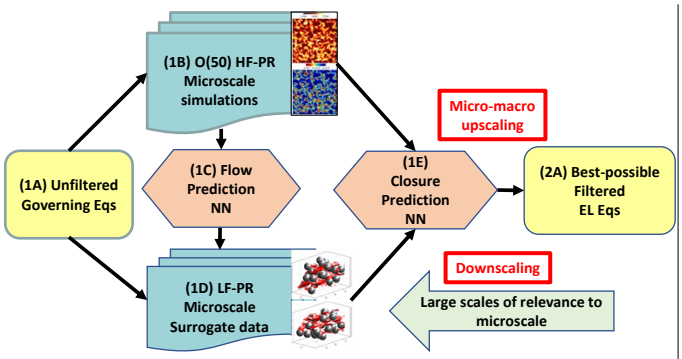

Operator Learning ML Model (OMLM)
Training/testing using HF-PR & LF-PR data

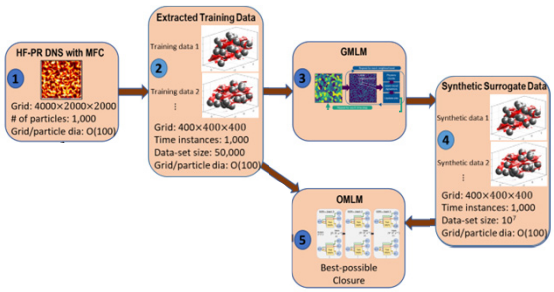

1: Perform the exascale HF-PR simulations

2: Curate to extract training data

3: Train the GLML (flow prediction)

4: Use HR-PR and LF-surrogate to develop OLML

CM³C

PR Simulation ⇒ ML ⇒ EL Simulation Pipeline

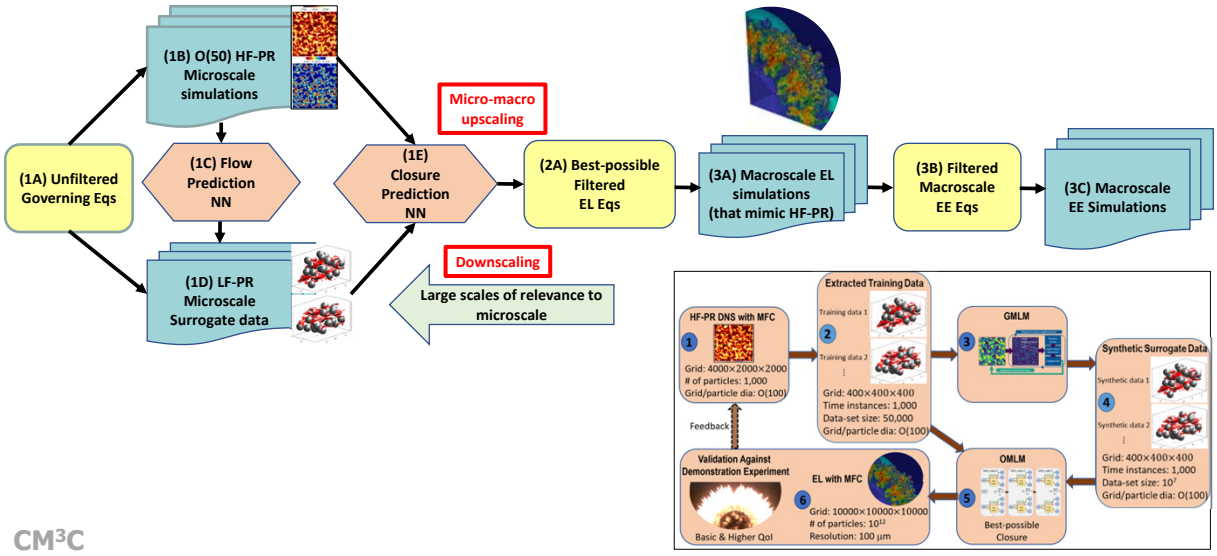

CM³C

Uncertainty Quantification & Propagation

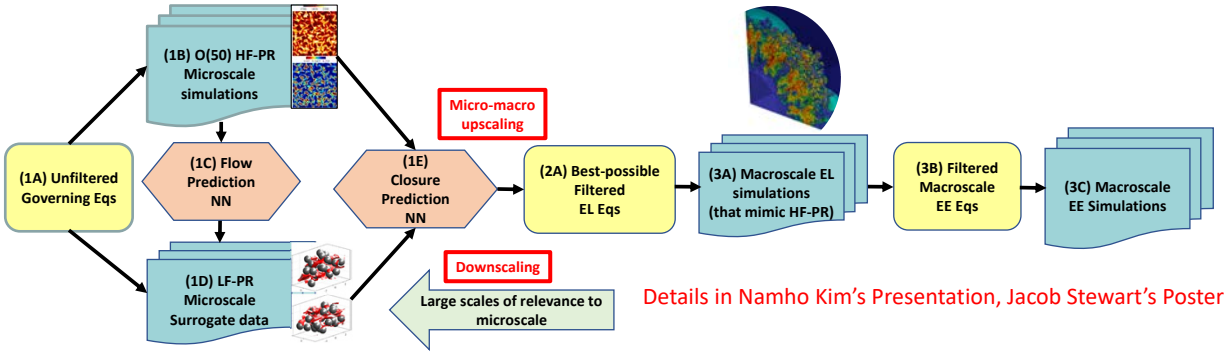

Details in Namho Kim's Presentation, Jacob Stewart's Poster

- Exascale HF-PR Simulations (1B)
  - Epistemic uncertainty from chemistry & radiation
  - Insufficient resolution at higher Re
  - Unrepresented physics
- LF-PR Surrogates (1D)
  - HF-PR uncertainties propagated through GMLM
  - Insufficient HF-PR data for training
  - Epistemic uncertainty from GMLM architecture limitations & training
- Macroscale EL Simulations (3A)
  - Aleatoric uncertainty from subgrid turbulence
  - Propagated uncertainty from the OMLM training
  - Insufficient OMLM training &architecture

CM³C

Validation-Quality Experimental Campaigns

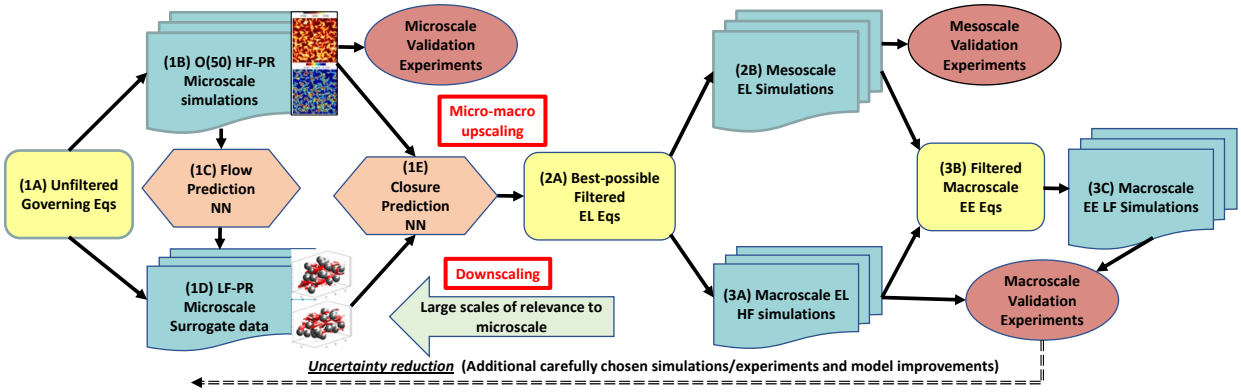

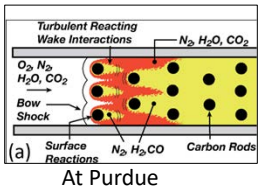

At Purdue

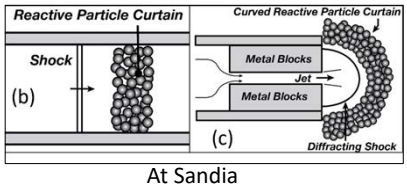

At Sandia

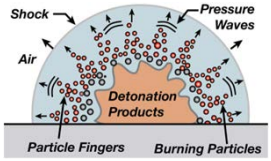

At Purdue, Eglin, Crane

CM³C

Envisioned Single Most Lasting Contribution After 5 Years

Revolutionary Modular Superposition (Compositional AI) framework that melds exascale simulations and "best-possible" ML closures to be used with unprecedented confidence in engineering analysis of multiphase flows

CM³C

Particle-Resolved Hero Simulation on 1024 GPU Nodes

- Hemispherical particle mound
- Compressed gas models detonation
  1. 100,000 resolved particles
  2. 20 cell diameters per particle
  3. Mach 25 hemispherical blast
  4. Ran for 8 hours on 1024 nodes of OLCF Frontier

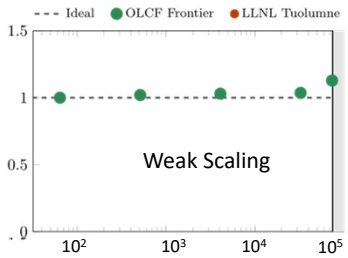

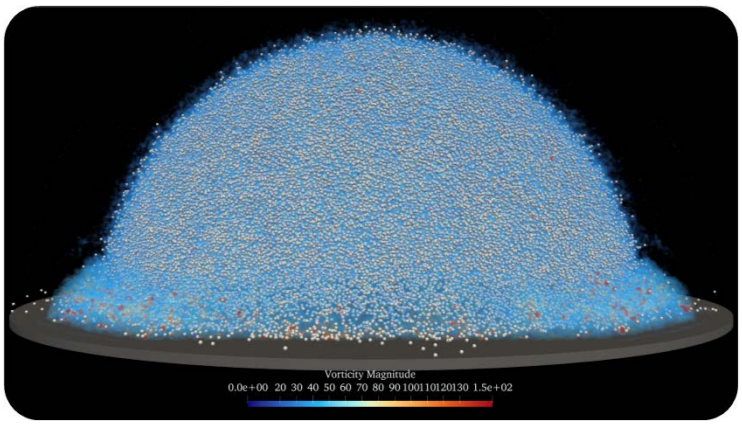

Details in Dan Vicker's Poster & Presentation

CM³C

EL Framework & Exascale Implementation

Magnesium particle shell
~ ½ million particles
Radius: 11.05 – 12.65 mm
100-micron diameter
20% peak volume fraction

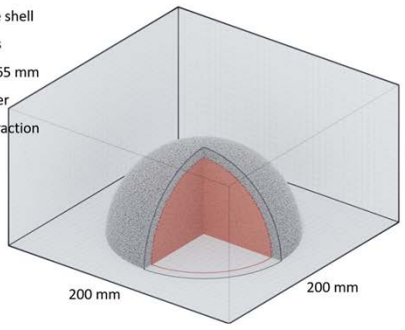

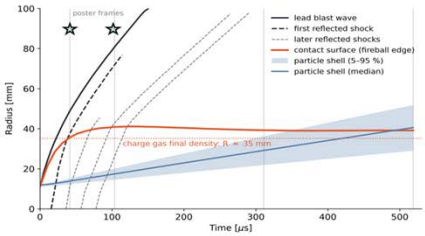

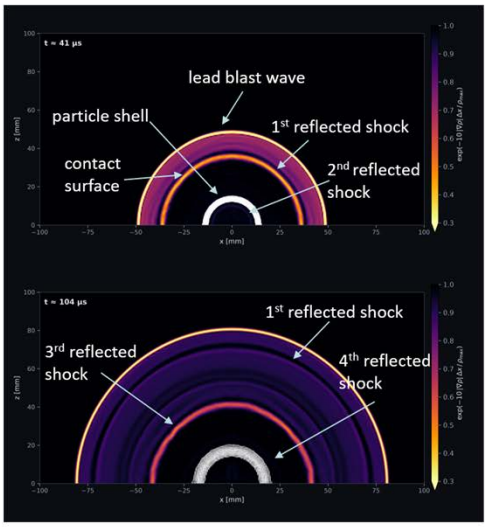

Details in Thierry Daoud's Poster & Presentation

Shock-Particle Interaction PR-DNS Data for ML Modeling

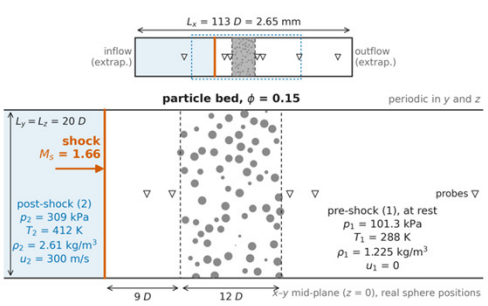

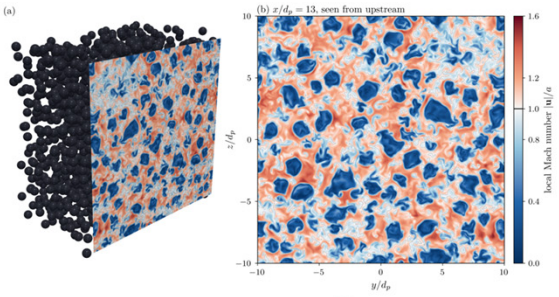

Details in Fahad Nabid's Poster

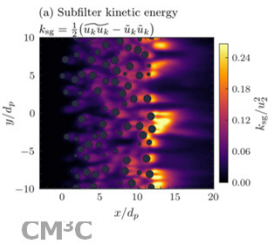

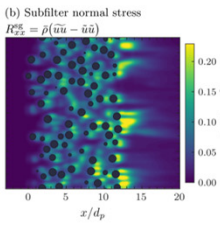

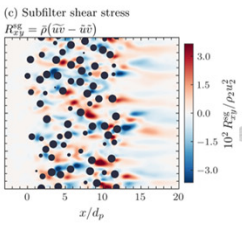

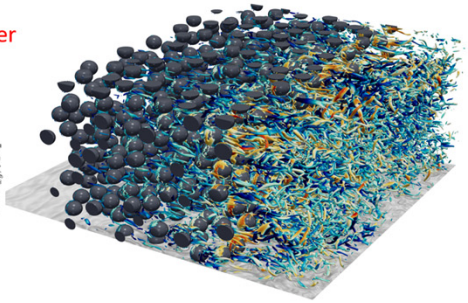

ML Modeling Pipeline – Flow Prediction

Nested ML training & Self-consistent update of Neighbor-influence

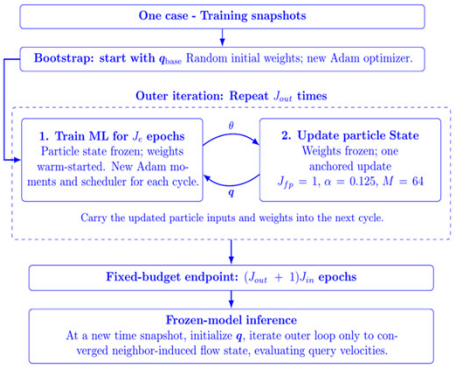

Details in Raul Valle's Talk & Aayushma Aashu's Poster

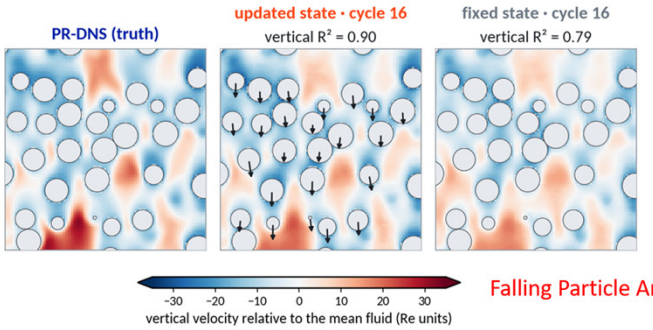

Falling Particle Array

Compressible Flow Past a Particle Array

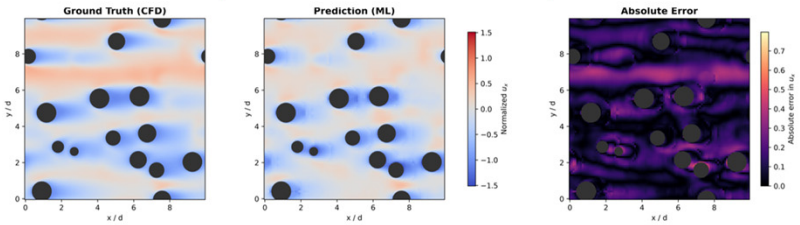

CM³C

Chemistry modeling in MFC

- Heterogenous surface reactions recently added to MFC for arbitrary combustion models
  1. MFC release V5.7.1 in Oct. 2026
  2. Can treat isothermal or energy flux balance conditions at particle surface
- Ex: carbon combustion
  1. Gas phase – GRI30 with 11 species and 33 reactions
  2. Surface reactions – 5 heterogeneous surface reactions (Chelliah *et. al.*, 1996)
- For magnesium, see poster by Chrissie Spangler and Gregory O'Brien

Details in T2. Ryan Houim, T3. Spencer Bryngelson, and T7. Davin Anderson

Details in posters by Davin Anderson, Isabella Rivera-Sigmann, Chrissie Spangler & Greg O'Brien (joint), Benjamin Wilfong

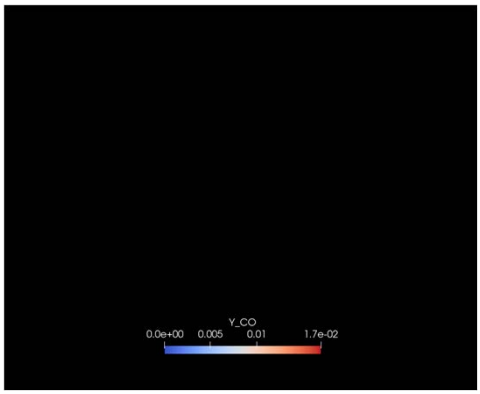

Isosurfaces of carbon mass fraction

Flow over a 3D sphere – Isabella Rivera-Sigmann
Computed on Tuolumne using MFC

CM³C

Carbon Rod Experiments and Simulations for Validation

- Carbon combustion models are being developed and validated against Purdue's carbon rod experiments
  1. Single rod case
  2. Multiple rod case
- 2D configuration makes it desirable to run many scenarios
- For more details, see
  1. T2. Ryan Houim
  2. T3. Spencer Bryngelson
  3. Posters: Benjamin Wilfong, Alex Millauer-Hsu

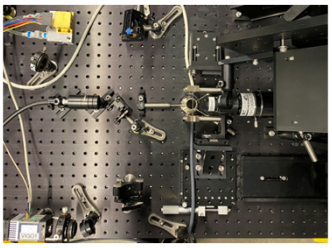

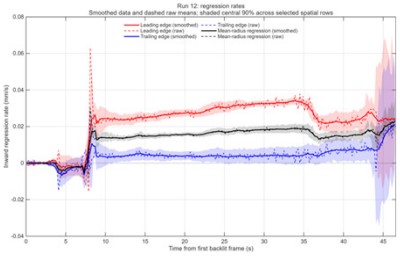

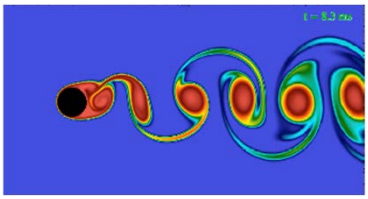

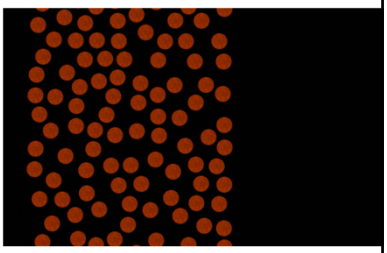

CM³C

Q3. Five-Year Prediction Plan

- Various aspects of 5-year plan are on target
  1. Simulations
  2. ML development
  3. Carbon rod experiments
  4. UQ microscale sims. and exps.
- MFC releases on regular basis
  1. Started: V5.0.0
  2. Now: V5.7.1
  3. Soon: V6.0.0
  4. Each PR requires extensive testing

	Y1	Y2	Y3	Y4	Y5		Y1	Y2	Y3	Y4	Y5
Simulations						**Experiments**					
HF-PR Micro (C rods)						Micro C rods					
HF-PR Micro (C & Mg)						Meso Campaigns					
EL Mesoscale						Macro (PU)					
EL Macroscale						Macro (other)					
ML development						**CS/Exascale**					
Generative ML						APU (El Capitan)					
Surrogate data						Dyn. Load Bal.					
OMLM deterministic						Integr. ML-closure					
Mori Zwanzig						Trace-based					
VVUQ						**Software**					
Relay validation						Software upgrades					
UQ micro						Regression testing					
UP across scales						Release - Minor					
UB and UR						Release - Major					

Table 1: Gantt chart for CM³C that shows the timeline of all the important activities of the center. Planned In progress Completed

See T2. Ryan Houim

CM³C

Key Physics – Yearly Breakdown Roadmap

- Several major objectives satisfied in Y1
  1. EOS
     a. MFC major refactoring to allow for Ideal, Stiffened, MG, JWL, BKW
  2. Chemical reactions
     a. MFC extended to include heterogeneous surface reactions
     b. Can handle 100's of reactions
  3. Euler-Lagrange
     a. Contact model (soft sphere, frictional)
     b. High-order interpolation and projection
     c. Periodic boundary condition
     d. State-of-the-art closure models
  4. Hero Simulations
     a. Modest in Y1
     b. Expected to be significant in Y2

	Y1	Y2	Y3	Y4	Y5
Gas-phase physics and MFC development					
Gas phase	LES	LES	LES	LES	LES
EOS	Ideal	JWL	BKW	BKW	BKW
Closure model	Standard Smag.	Dynamic Smag.	Dynamic Smag.	ML-based	ML-based
Chem. reactions	$\mathcal{O}(1)$	$\mathcal{O}(10)$	$\mathcal{O}(10)$	$\mathcal{O}(100)$	$\mathcal{O}(100)$
Wall effects	Inviscid	Inviscid	Inviscid	Inviscid	Viscous
Grid size	10^{12}	10^{12}	10^{12}	10^{12}	10^{13}
Euler–Lagrange MFC developments and details					
Dispersion (EL)	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg
Concentration	Dilute	Dilute	Dense	Dense	Dense
Contact model	None	Soft sphere	Frictional soft sphere	Frictional soft sphere	Frictional soft sphere
Force model	Standard	Standard	ML-based	ML-based	ML-based
Particles	10^{10}	10^{11}	10^{12}	10^{13}	10^{13}
Load balancing	Simple	Bin-based	Bin-based	Trace-based	Trace-based
Hero Sims (EL)	2	2	4	4	4
Euler–Euler MFC developments and details					
Dispersion (EE)	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg
Distribution	Monodisperse	Polydisperse	Polydisperse	Polydisperse	Polydisperse
Grid size	10^{10}	10^{10}	10^{10}	10^{10}	10^{10}
Grid refinement	Static	Static	AMR	AMR	AMR
Concentration	Dilute–dense	Dilute–dense	Dilute–dense	Dilute–dense	Dilute–dense
Contact model	Granular	Granular	Granular	ML-based	ML-based
Force model	Standard	Standard	ML-based	ML-based	ML-based
EE Sims.	50	100	300	600	1000
Exascale plan					
Architecture	GPU (Lassen) APU (El Capitan)	APU (El Capitan)	APU (El Capitan)	APU (El Capitan)	APU (El Capitan)
Machine utilization	~100%	~100%	~100%	~100%	~100%

See T2. Ryan Houim

ASC TriLab Supercomputing Resources

- Tuolumne, Dane, Frontier
- Users
  1. UF – 16 (Phys, ML, UQ, Group Comb. teams)
  2. UTA/GT – 3
  3. Purdue – pending after AST
- Usage
  1. Florida
     a. Y1: mainly test runs for PRS, EL, and carbon chemistry developments
     b. Y2: larger production runs and 3D carbon combustion
  2. UTA/GT – Development, Scaling and larger physics-based runs

Florida Usage

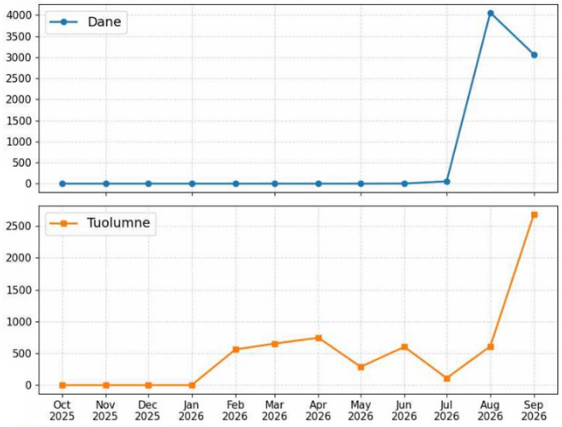

Center Organization Chart

- Teams
  1. Phys-S: Microscale PR-DNS
  2. Phys-S: Meso-Macroscale EL
  3. CS-Exa: MFC Comput. Science
  4. Phys-E: Experiments
  5. ML: Machine Learning
  6. VVUQ: Relay Validation/Workflow
  7. Group Combustion
- Members fit into multiple teams

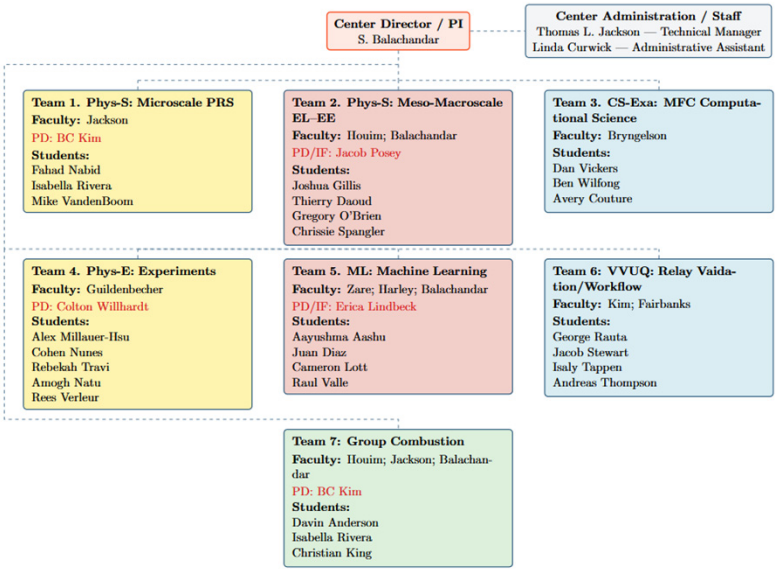

CM³C

Integration of Tasks and Teams

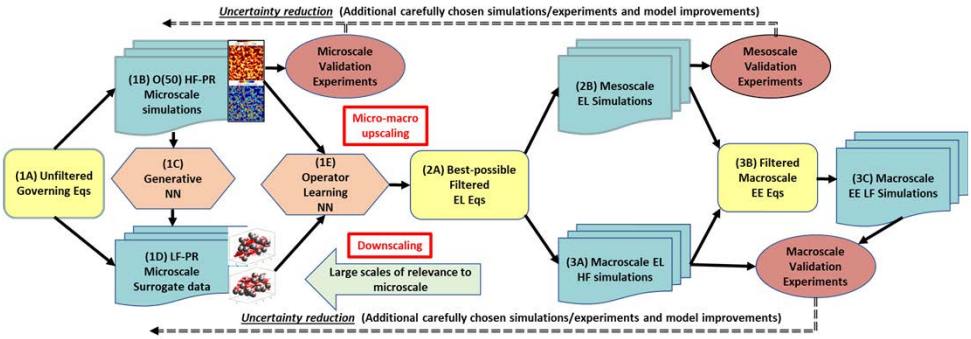

PR/Comb	ML	EL/EE/Comb	MFC	Exp	UQ/Workflow
• BC Kim	• Erica Lindbeck	• Jacob Posey	• Benjamin Wilfong	• Colton Willhardt	• Jacob Stewart
• Mike VandenBoom	• Isaly Tappan	• Joshua Gillis	• Daniel Vickers	• Alex Millauer-Hsu	• George Rauta
• Fahad Nabid	• Cameron Lott	• Gregory O'Brien	• Thierry Daoud	• Cohen Nunes	• Andreas Thompson
• Davin Anderson	• Aayushma Aashu	• Chrissie Spangler	• Avery Couture	• Rebekah Travis	
• Isabella Rivera-Sigmann	• Juan Diaz				
	• Christian King				
	• Julian Reyes				
	• Raul Valle				

Postdoc
Undergrad

CM³C

Student Internship - Completed

Name	Laboratory	Mentor	Year/ Length (wks)	Research
Thierry Daoud	LLNL	Georges Akiki	2026/10	Pseudo-turbulence modeling connected with multiphase flows
Cohen Nunes	Sandia	Michael Kaneshige	2026/10	Experimental Detonation Sciences
Benjamin Wilfong	LLNL	Britton Olson	2026/10	High speed multiphase CFD, models and numerics for strong density differences and impact.

Q6. Student Internship - Planned

Name	Position	Degree/Date	Advisor	Internship
Davin Anderson	PhD (UF)	MAE (2029)	Balachandar	Summer 2027
Joshua Gillis	PhD (UF)	MAE (2029)	Balachandar	Summer 2027
Isabella Rivera-Sigmann	PhD (UF)	MAE (2029)	Balachandar	Summer 2027
Chrissie Spangler	PhD (Purdue)	ME (2029)	Houim	Summer 2027
Gregory O'Brien	PhD (Purdue)	ME (2029)	Houim	Summer 2027
Alex Millauer-Hsu	PhD (Purdue)	ME (2030)	Guildenbecher	Summer 2028
Daniel J. Vickers	PhD (GT)	CSE (2030)	Bryngelson	Summer 2027
George Rauta	PhD (UF)	ME (2029)	Fairbanks	Summer 2027
Raul Valle	PhD (UF)	ECE (2030)	Harley	Summer 2028
Isaly Tappan	PhD (UF)	ECE (2027)	Harley	Summer 2027
Juan Diaz	PhD (UF)	ECE (2029)	Zare	Summer 2027

Integration of Tasks and Teams

Tasks

- T1: Exascale simulations and experiments of demonstration problem (TM)
- T2: Integration of ML for scale-bridging (PD2)
- T3: UQ-driven micro and mesoscale investigations (PD3)
- T4: Multi-fidelity data/workflow (PD2)
- T5: Load balancing for exascale performance (PD1)

	Phys-S&E Director, TM Guildenbecher Houim Students	VVUQ/WF Kim, Fairbanks Students	ML Zare, Harley Director Students	CD-N Houim Bryngelson Director Students	CS-Exa Bryngelson Harley Students
Task 1 Lead: TM	X	X		X	X
Task 2 Lead: PD2	X	X	X		X
Task 3 Lead: PD3	X	X	X	X	
Task 4 Lead: PD2	X		X		X
Task 5 Lead: PD1				X	X

The Center will be organized by physics-based tasks and cross-cutting teams, rather than by faculty and their research groups

1. Weekly student/staff-led team interactions
   a. See the "big picture", Own their problem
   b. Become invested in the success of the center!
2. All staff and large number of graduate students are co-located (*4th floor of NEB*)
3. Team meetings held in Student Village conference room, 4th floor NEB

CM³C

CM³C

Questions?

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

UF UNIVERSITY of FLORIDA

Challenges Resolved and Challenges Ahead

- Reduced order chemistry has too much uncertainty
  - MFC can do handled detailed chemistry
- Radiation between particles (or rods) is non negligible
  - Yes, simple radiation will need to be implemented
- Load-balanced exascale EL implementation in MFC
  - MFC is very versatile – we do not see this as a major issue
- Not enough PR data for Generative ML model
  - Running big/bundled simulations with MFC is easy
- Initial compaction phase introduces complex unaccounted physics
- Superposition model fails for finite Re combustion
- Generative + Operator Learning models introduce too much uncertainty
- Availability of computing resources

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

T2. Annual Prediction Status --- Integration/workflow of required physics modules/codes and scaling

Co-Principal Investigator: Ryan W. Houim

School of Mechanical Engineering, Purdue University

School of Aeronautics and Astronautics, Purdue University

NNSA National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

PURDUE UNIVERSITY

Annual Prediction Questions

- **Q1.** The goals of your first full system demonstration(s), either in year 1 (optional) or year 2 (mandatory). How does it contribute to your ultimate year-five predictive simulation?
- **Q2.** If you were able to make a first prediction, what were the results? If not, what progress have you made towards a first prediction, and what remains to be done? In either case, what insights into predictiveness ("a-ha moments") were obtained?
- **Q3.** Relative to the year-five prediction, what can and cannot be computed and predicted at this point?
- **Q4.** What are the key physics components still missing from your simulation as determined by your roadmap and UQ process, and when do you expect to incorporate them?
- **Q5.** Describe the scaling and performance obtained, including a discussion of any limitations encountered. (See T3 by Spencer Bryngelson)
- **Q6.** How have you verified and validated your simulations?

CM³C Reactive Flow Physics Simulation Team


S. Balachandar *(Florida)*


Spencer Bryngelson *(UT Austin)*


Ryan Houim *(Purdue)*


Thomas Jackson *(Florida)*


ST Thakur *(Florida)*


Jacob Posey *(Purdue)*


BC Kim *(Florida)*


Davin Anderson *(Florida)*


Avery Couture *(Florida)*

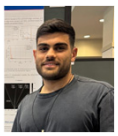
Thierry Daoud *(Florida)*


Joshua Gillis (Florida)


Christian King *(Florida)*


Cameron Lott *(Florida)*


Fahad Nabid *(Florida)*


Gregory O'Brian *(Purdue)*


Isabella Rivera *(Florida)*


Chrissie Spangler *(Purdue)*


Mike VandenBoom *(Florida)*


Daniel Vickers *(GT)*


Ben Wilfong *(GT)*

CM³C

Integration of Physics and Simulations

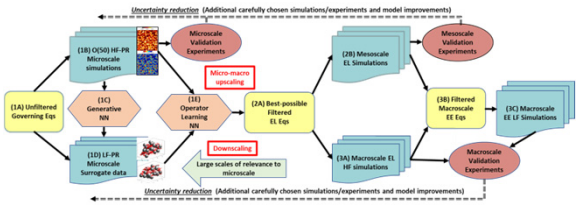

Integration Roadmap

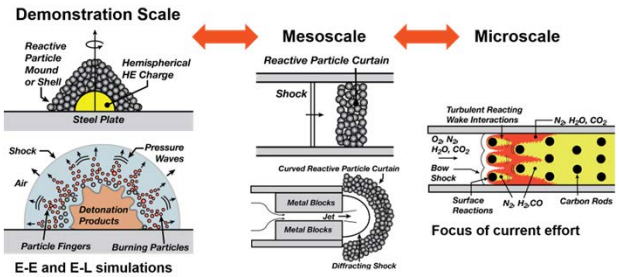

E-E and E-L simulations performed to assist experimental team

Discussion with SNL personnel on proposal development

Physics & Simulations

First part of talk

Second part of talk

CM³C

PURDUE UNIVERSITY

Q3. Five-Year Prediction Plan

- Various aspects of 5-year plan are on target
  1. Simulations
  2. ML development
  3. Carbon rod experiments
  4. UQ microscale sims. and exps.
- MFC releases on regular basis
  1. Started: V5.0.0
  2. Now: V5.7.1
  3. Soon: V6.0.0
  4. Each PR requires extensive testing

	Y1	Y2	Y3	Y4	Y5		Y1	Y2	Y3	Y4	Y5
Simulations						**Experiments**					
HF-PR Micro (C rods)						Micro C rods					
HF-PR Micro (C & Mg)						Meso Campaigns					
EL Mesoscale						Macro (PU)					
EL Macroscale						Macro (other)					
ML development						**CS/Exascale**					
Generative ML						APU (El Capitan)					
Surrogate data						Dyn. Load Bal.					
OMLM deterministic						Integr. ML-closure					
Mori Zwanzig						Trace-based					
VVUQ						**Software**					
Relay validation						Software upgrades					
UQ micro						Regression testing					
UP across scales						Release - Minor					
UB and UR						Release - Major					

Table 1: Gantt chart for CM^{3}C that shows the timeline of all the important activities of the center.
Planned In progress Completed

CM³C

PURDUE UNIVERSITY

Q4. Key Physics – Yearly Breakdown Roadmap

- Several major objectives satisfied in Y1
  1. EOS
     a. MFC major refactoring to allow for Ideal, Stiffened, MG, JWL, BKW
  2. Chemical reactions
     a. MFC extended to include heterogeneous surface reactions
     b. Can handle 100's of reactions
  3. Euler-Lagrange
     a. Contact model (soft sphere, frictional)
     b. High-order interpolation and projection
     c. Periodic boundary condition
     d. State-of-the-art closure models
  4. Hero Simulations
     a. Modest in Y1
     b. Expected to be significant in Y2

	Y1	Y2	Y3	Y4	Y5
	Gas-phase physics and MFC development				
Gas phase	LES	LES	LES	LES	LES
EOS	Ideal	JWL	BKW	BKW	BKW
Closure model	Standard Smag.	Dynamic Smag.	Dynamic Smag.	ML-based	ML-based
Chem. reactions	$\mathcal{O}(1)$	$\mathcal{O}(10)$	$\mathcal{O}(10)$	$\mathcal{O}(100)$	$\mathcal{O}(100)$
Wall effects	Inviscid	Inviscid	Inviscid	Inviscid	Viscous
Grid size	10^{12}	10^{12}	10^{12}	10^{12}	10^{13}
	Euler–Lagrange MFC developments and details				
Dispersion (EL)	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg
Concentration	Dilute	Dilute	Dense	Dense	Dense
Contact model	None	Soft sphere	Frictional soft sphere	Frictional soft sphere	Frictional soft sphere
Force model	Standard	Standard	ML-based	ML-based	ML-based
Particles	10^{10}	10^{11}	10^{12}	10^{13}	10^{13}
Load balancing	Simple	Bin-based	Bin-based	Trace-based	Trace-based
Hero Sims (EL)	2	2	4	4	4
	Euler–Euler MFC developments and details				
Dispersion (EE)	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg
Distribution	Monodisperse	Polydisperse	Polydisperse	Polydisperse	Polydisperse
Grid size	10^{10}	10^{10}	10^{10}	10^{10}	10^{10}
Grid refinement	Static	Static	AMR	AMR	AMR
Concentration	Dilute-dense	Dilute-dense	Dilute-dense	Dilute-dense	Dilute-dense
Contact model	Granular	Granular	Granular	ML-based	ML-based
Force model	Standard	Standard	ML-based	ML-based	ML-based
EE Sims.	50	100	300	600	1000
	Exascale plan				
Architecture	GPU (Lassen) APU (El Capitan)	APU (El Capitan)	APU (El Capitan)	APU (El Capitan)	APU (El Capitan)
Machine utilization	~100%	~100%	~100%	~100%	~100%

CM³C

Q1-3. Goals of Simulation Team

- Over-arching Goal
  1. Gain physical insight into Group Combustion
  2. Develop appropriate models and numerical tools for droplet combustion
  3. Develop ML models for Euler-Lagrange simulations at the demonstration scale
- Carry out simulations at all three scales
  1. Demonstration problem
  2. Mesoscale problem
  3. Microscale
  4. Both inert and reacting particles
- Carry out code verification and validation
- Provide extensive data base for ML team

Particle Resolved Simulations

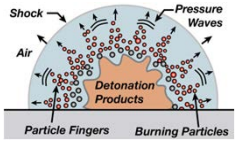

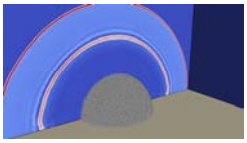

Particle Resolved Simulations

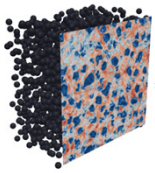

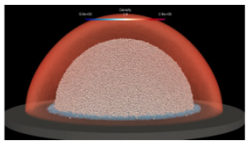

Group Combustion using PRS

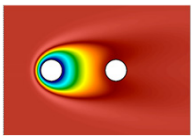

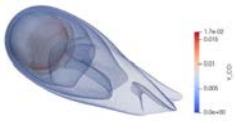

CM³C

Q1-3. Demonstration Problem

- Y1 progress
  - 6 force/closure models have been added
  - Two-way coupling algorithms (interpolation and projection)
  - 20 regression tests, with several example cases
- Y1 largest simulation
  - 70 million individually tracked Lagrangian particles
  - 10 nodes on LLNL Tuolumne, 40 AMD MI300A GPUs
  - 32 million fluid grid cells
- Y2 effort
  - Addition of 4 more closure models
  - JWL EOS support for Lagrangian particles
  - Performance and memory optimization up to machine scale
  - Full scale simulation (2)

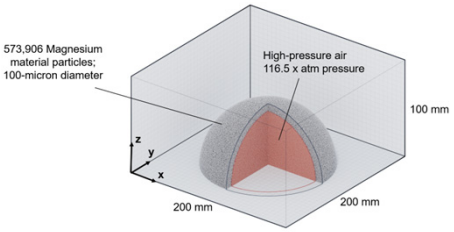

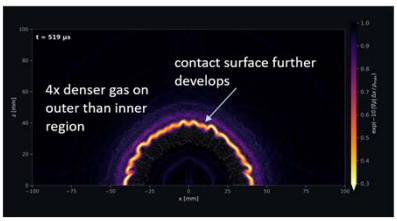

- See T6 by Thierry Daoud
- See posters by Joshua Gillis, Thierry Daoud, Ben Wilfong, Avery Couture

CM³C

Q1-3. EE Macroscale Simulations for Experimental Design

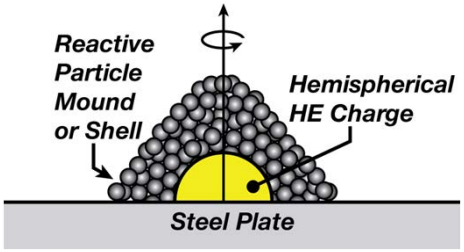

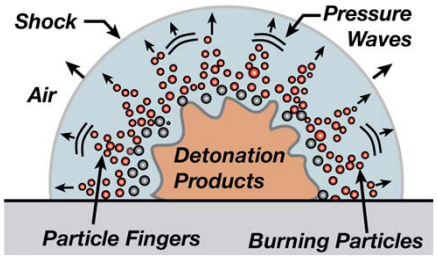

- Simulate explosive dispersal, ignition, and combustion of Mg particles by PETN with the HyBurn code to aid design of experiments at Purdue
- Eulerian model used for the multiphase mixture and programmed burn for the explosive
- Simplified approach for thermodynamics and reactions based on our previous work with Al particles

Houim & Oran, J. Fluid Mech. (2016)
Houim, Shock Waves, (2021)
Posey *et al.*, Phys. Fluids (2021)

CM³C

Q1-3. Three Phase EE Multiphase Flow Model

- Used with Guildenbecher group to aid in experimental design and implement new Mg models as they are developed.
- The models and algorithms developed in-house[1-3]
- The three-phases are: gas, explosive, and particles (Eulerian)
- A diffuse interface method (DIM) is used for reactive gas and high-explosive (HE)
  - Novel detonation burn model converts HE to products at the CJ composition[1,3]
- An Eulerian multiphase model is used for the particles[2]
  - Accounts for phase change, drag, heat transfer, collisions, and compaction
  - Particles burn via an empirical burning model
- Turbulence modeled using ILES[4]

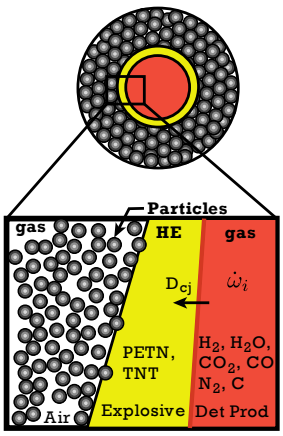

[1]Houim, *Shock Waves* (2021)
[2]Houim, Oran, *J. Fluid Mech.* (2016)
[3]Egeln, *et al.*, *Phys. Fluids*, (2023)
[4]Thornber, *et al.*, *J. Comp. Phys*, (2008)

CM³C

Q1-3. EE Macroscale Runs – Preliminary Results

- 6 mm PETN charge with 3 mm-thick shell of 10 μm-diameter Mg particles at 50% packing
  - Use simple Mg burning rate correlation for combustion.
- Results show the show the dispersal, ignition, and combustion of the Mg particles.

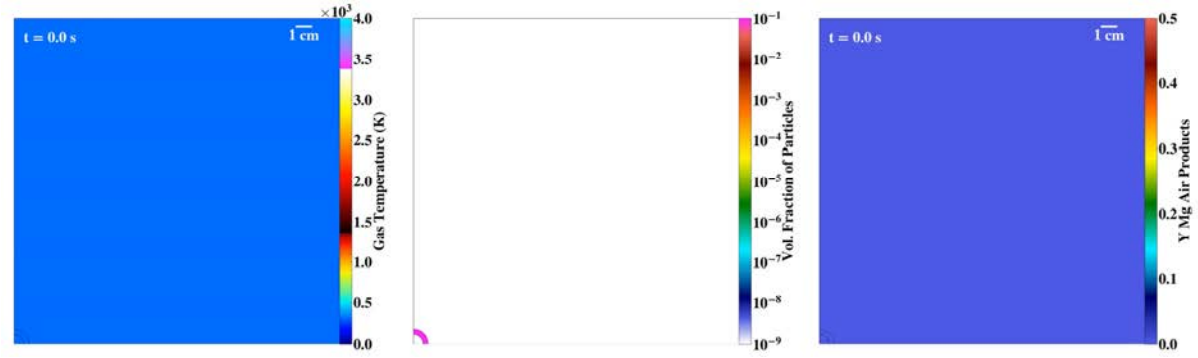

CM³C

Q1-3. Mesoscale Simulations

- Mesoscale problems
  1. Planar curtain
     a. Particle-resolved simulations (PRS)
     b. Euler-Lagrange simulations (EL)
     c. Verification of EL using PRS
  2. Hemispherical blast
     a. PRS
     b. Euler-Lagrange
  3. For work on PRS, see T6 by Daniel Vickers
- One goal is to understand and develop models for all closure terms that can be used at meso and macroscale simulations
  1. This information will also be used by the ML Team to develop comprehensive operators for the closure terms
  2. *See next slide*

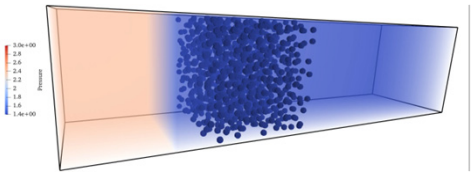

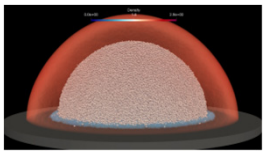

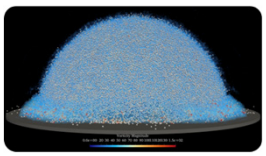

- See posters by Fahad Nabid, Cameron Lott, Daniel Vickers, Mike VandenBoom

CM³C

Q1-3. Mesoscale Simulations – Closure Terms (non-reacting)

- The exact phase averaged compressible equations, with no small perturbation assumption:

$$\frac{\partial}{\partial t}\left(\phi_g\,\overline{\rho}_g\right)+\nabla\cdot\left(\phi_g\,\overline{\rho}_g\,\widetilde{\mathbf{u}}_g\right)=0$$

$$\frac{\partial}{\partial t}\left(\phi_g\,\overline{\rho}_g\widetilde{\mathbf{u}}_g\right)+\nabla\cdot\left(\phi_g\,\overline{\rho}_g\widetilde{\mathbf{u}}_g\otimes\widetilde{\mathbf{u}}_g\right)=-\phi_g\nabla\overline{p}_g+\boxed{\phi_g\nabla\cdot\overline{\boldsymbol{\tau}}_g}-\boxed{\nabla\cdot\left(\phi_g\,\mathbf{R}_{sg}\right)}-\boxed{\mathbf{F}'_g}$$

$$\frac{\partial}{\partial t}\left(\phi_g\,\overline{\rho}_g\widetilde{E}_g\right)+\nabla\cdot\left(\phi_g\,\overline{\rho}_g\widetilde{\mathbf{u}}_g\widetilde{E}_g\right)=-\nabla\cdot\left(\phi_g\,\overline{p}_g\widetilde{\mathbf{u}}_g\right)+\boxed{\nabla\cdot\left(\phi_g\,\overline{\boldsymbol{\tau}}_g\cdot\widetilde{\mathbf{u}}_g\right)}$$

$$-\overline{p}_g\nabla\cdot\left(\phi_p\mathbf{u}_p\right)+\boxed{\overline{\boldsymbol{\tau}}_g:\nabla\left(\phi_p\mathbf{u}_p\right)-\phi_g\nabla\cdot\overline{\mathbf{q}}_g}-\boxed{\left(W'_q+Q'_q\right)}-\boxed{\nabla\cdot\left(\phi_g\,\mathbf{Q}_{sg}\right)}$$

Viscous/Diffusive | Surface Integrals | Closure Terms

- Filtered quantities to be computed, needed to develop ML model:

- $\mathbf{R}_{sg}=\overline{\rho\mathbf{u}\otimes\mathbf{u}}-\frac{1}{\overline{\rho}}\overline{\rho\mathbf{u}}\otimes\overline{\rho\mathbf{u}}$
- $\mathbf{Q}_T=\gamma\,C_v\left(\overline{\rho\mathbf{u}T}-\widetilde{T}\cdot\overline{\rho\mathbf{u}}\right)$,
- $\mathbf{E}_{ku}=\frac{1}{2}\left[\overline{\rho\mathbf{u}(\mathbf{u}\cdot\mathbf{u})}-\frac{\overline{\rho\mathbf{u}}}{\overline{\rho}}\overline{\rho(\mathbf{u}\cdot\mathbf{u})}\right]$
- $\mathbf{W}_{\tau u}=\frac{1}{\overline{\rho}}\left(\overline{\rho\boldsymbol{\tau}\cdot\mathbf{u}}-\overline{\boldsymbol{\tau}}\cdot\overline{\rho\mathbf{u}}\right)$
- $\mathbf{R}^g_{\mu,sg}=\overline{\boldsymbol{\tau}}-\frac{2}{3}\mu_g\left(\nabla\cdot\overline{\mathbf{u}}_m\right)\mathbf{I}+\mu_g\left[\left(\nabla\overline{\mathbf{u}}_m\right)+\left(\nabla\overline{\mathbf{u}}_m\right)^{\top}\right]$
- $\mathbf{R}^g_{\lambda,sg}=\overline{\mathbf{q}}-\lambda_g\nabla(\widetilde{T})$

- Filtered quantities needed from PR simulations can be sorted into scalar, vector, and second-rank tensor:

- **Scalar:** $\overline{\phi}_p$, $\overline{\rho}$, $\overline{\rho(\mathbf{u}\cdot\mathbf{u})}$
- **Vector:** $\overline{\mathbf{u}}_p$, $\overline{\rho\mathbf{u}}$, $\overline{\rho\mathbf{u}T}$, $\overline{\rho\mathbf{u}(\mathbf{u}\cdot\mathbf{u})}$, $\overline{\rho\boldsymbol{\tau}\cdot\mathbf{u}}$, $\overline{\mathbf{q}}$
- **Tensor:** $\overline{\rho\mathbf{u}\otimes\mathbf{u}}$, $\overline{\boldsymbol{\tau}}$

CM³C

Q1-3. Mesoscale Simulations – Closure Terms

- Example from PRS showing sub-filtered kinetic energy, normal stress, shear stress, enthalpy flux, kinetic-energy flux, and viscous work

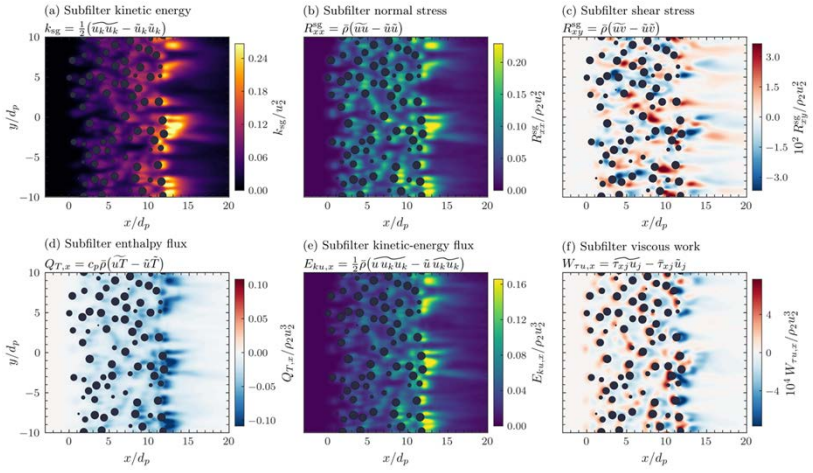

CM³C

$\overline{(\cdot)}$: phase-weighted Gaussian filter, $\sigma=d_p/4$; $\bar{\rho}=\overline{\rho}$, $\tilde{\phi}=\overline{\rho\phi}/\bar{\rho}$; $\bar{\tau}_{ij}$: filtered viscous stress; $u=u_x$, $v=u_y$

Q1-3. Microscale

- Microscale problems
  1. Carbon rod
  2. Validation with Purdue experiments
  3. Cross code verification – Hyburn and MFC (physics and implementation)
  4. Multiple carbon cylinders and spheres for group combustion – early stages of development



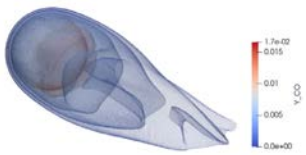

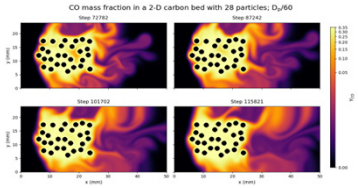

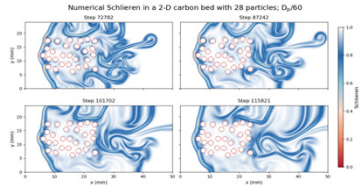

CM³C • See posters by Davin Anderson, Isabella Rivera-Sigmann, Christian King, Gregory O'Brian, Chrissie Spanger

Q1-3. Carbon Rod Model: Goals

- Develop and reduce the interface-resolved carbon burning model
  - Same model for particles and rods: just different geometry
  - Help with chemical mechanism reduction
  - Compare oxidation in wet and dry air
- Validate burning against experimental measurements
  - Makino *et al., Combust Flame,* 2003
  - Guildenbecher group at Purdue

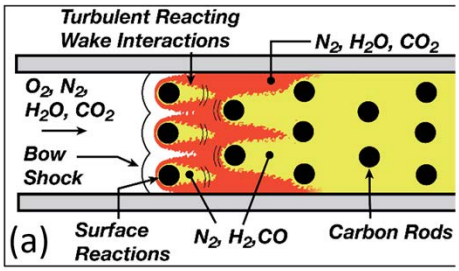

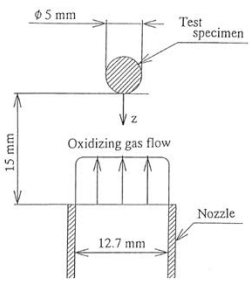

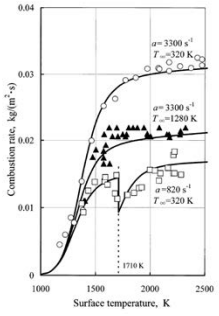

CM³C

Q1-3. Microscale Carbon Rod/Particle: Governing Equations

- Level Set/IBM used to separate gas and solid (carbon) regions
  - regression implemented
- Solid Region: heat equation

$$\frac{\partial T}{\partial t} + \mathbf{v} \cdot \nabla T = \frac{1}{\rho C} \nabla \cdot \lambda \nabla T + \frac{\dot{q}_{\text{gen}}}{\rho C}$$

- Gas Region: Reactive N-S, with Mixture-averaged molecular transport

$$\frac{\partial \rho Y_i}{\partial t} + \nabla \cdot [\rho Y_i (\mathbf{v} + \mathbf{V}_i)] = \dot{\omega}_i,$$

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{v}) + \nabla p = \nabla \cdot \sigma,$$

$$\frac{\partial \rho E}{\partial t} + \nabla \cdot [\mathbf{v}(\rho E + p)] = -\nabla \cdot \mathbf{q} + \nabla \cdot (\sigma \cdot \mathbf{v}),$$

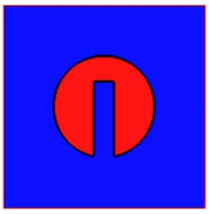

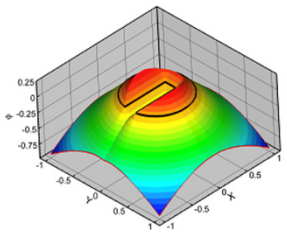

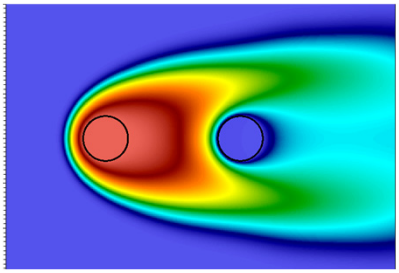

CM³C

Q1-3. Microscale Surface Boundary Conditions

- *Mass Balance*: Gives gas-phase Stefan ("blowing") velocity jump at surface.

$$u_g = u_c + \dot{m}_c'' \left(\frac{1}{\rho_g} - \frac{1}{\rho_c} \right)$$

- *Species Mass Balance*: The mass production rate due to surface reaction is balanced with species mass advection and diffusion

$$\dot{\omega}_i'' = \rho_g Y_i (u_g + V_i)$$

- *Energy Balance*

$$-\lambda_g \frac{\partial T}{\partial n} + \sum \rho Y_{s,i} h_i(T_s)(u_s + V_i) +$$

$$\dot{q}_{\text{rad,out}} - \dot{q}_{\text{rad,in}} = -\lambda_s \frac{\partial T}{\partial n} + \dot{m}_c'' h_c(T_s)$$

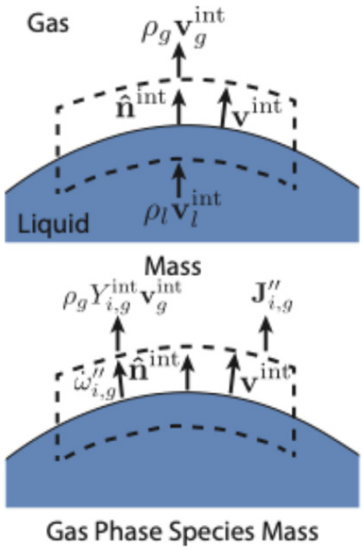

CM³C

Q1-3. Carbon Surface and Gas Chemistry

- Originally from Bradley *et al*. ProCI, 1985.
- Accounts for wet and dry oxidation of carbon
- Involves 5 reactions and 8 gas species
- Gas Chemistry: Syngas mechanism w/12 species and 36 rxns. Li *et al.* ProcCI, 2014.

$$C(s) + \frac{1}{2}O_2 \rightarrow CO, \quad (15)$$

$$C(s) + CO_2 \rightarrow 2CO, \quad (16)$$

$$C(s) + H_2O \rightarrow CO + H_2, \quad (17)$$

$$C(s) + OH \rightarrow CO + H, \quad (18)$$

$$C(s) + O \rightarrow CO, \quad (19)$$

Table 1
Kinetic parameters for heterogeneous reactions Eqs. (15)–(19) taken from Bradley et al. [49].

r	B_r	E_r [kJ/mol]	A_r	R_r [kg/(m²s)]
15*a*	0	125.6	2400 [kg/(m² s atm)]	$R_{15} = \frac{K_{15a}P_{O_2}\xi}{1+K_{15b}P_{O_2}} + K_{15c}P_{O_2}(1-\xi)$
15*b*	0	−17.17	21.3 [1/atm]	
15*c*	0	63.64	0.535 [kg/(m² s atm)]	$\xi = 1/\left(1+\frac{K_{15d}}{K_{15c}P_{O_2}}\right)$
15*d*	0	406.1	18.1×10^6 [kg/(m² s)]	
16	0	285	9×10^3 [kg/(m² s atm$^{0.5}$)]	$R_{16} = K_{16}P_{CO_2}^{0.5}$
17	0	288	4.8×10^5 [kg/(m² s atm$^{0.5}$)]	$R_{17} = K_{17}P_{H_2O}^{0.5}$
18	−0.5	0	361 [kg K$^{0.5}$/(m² s atm)]	$R_{18} = K_{18}P_{OH}$
19	−0.5	0	665.5 [kg K$^{0.5}$/(m² s atm)]	$R_{19} = K_{19}P_O$

Q1-3. Reactive Wall Model

- Assumptions:
  - Quasi-steady
  - Fickian Diffusion
  - Constant ρD
  - No gas-phase reactions between surface and IP

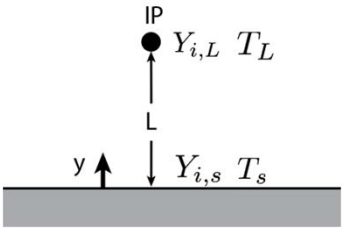

Surface species mass balance

$$\dot{\omega}_i'' = \dot{m}_c'' Y_{i,s} - \rho \mathcal{D} \left.\frac{\partial Y_i}{\partial y}\right|_s \approx \dot{m}_c'' Y_{i,s} - \rho \mathcal{D} \frac{Y_{i,L} - Y_{i,s}}{L}$$

Known mass fraction at IP determined from interpolation $Y_i(L) = Y_{i,L}$

Surface reaction model $\dot{\omega}_i'' = f(Y_s, T_s, p_s, \ldots) \quad \dot{m}_c'' = \sum \dot{\omega}_i''$

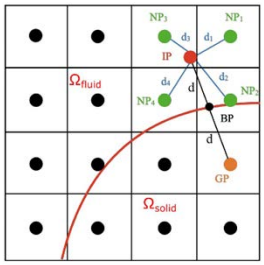

CM³C

Q1-3. Low-Mach and Diffusion Speedups

- Implemented Iterative Dynamic Chemical Stiffness Removal (IDSCR)
  - ~1.7x speed up for *entire* rod sim over VODE.
- Implemented Reduced Speed of Sound Technique (RSST) to mitigate acoustic CFL limit.
  - Adapted to combustion from low-Mach astrophysical simulation techniques.
  - Scale ∇p by $\beta = \left(\frac{M_{\text{ref}}}{M_{\text{target}}}\right)^2$, which gives $c_{\text{num}} = \sqrt{\beta} C_{\text{real}}$
- Runge-Kutta-Legendre (RKL) time-stepping for diffusion.
  - Explicit super time-stepping method for parabolic PDEs
- All together: ~23x speed up over time-marching fully compressible eqns.

Stoichiometric H_2-O_2 Flame

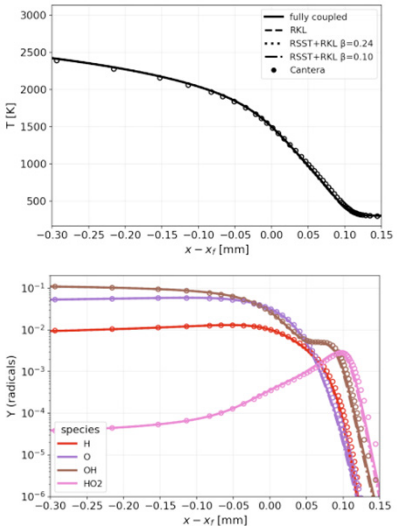

CM³C

Q1-3. Preliminary Comparison to Makino et al. Data

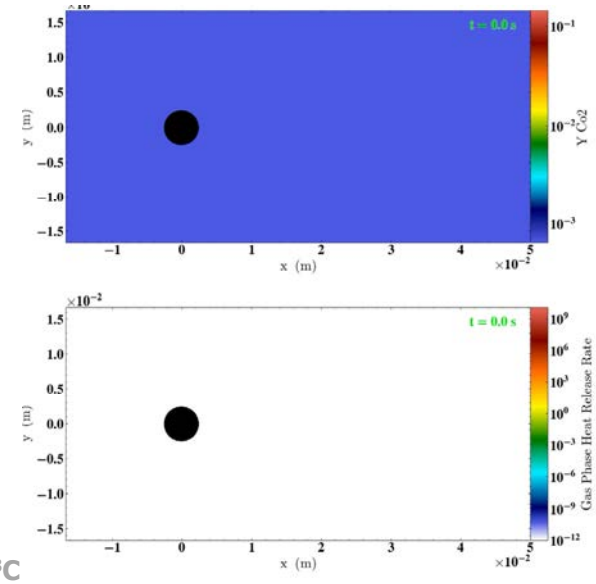

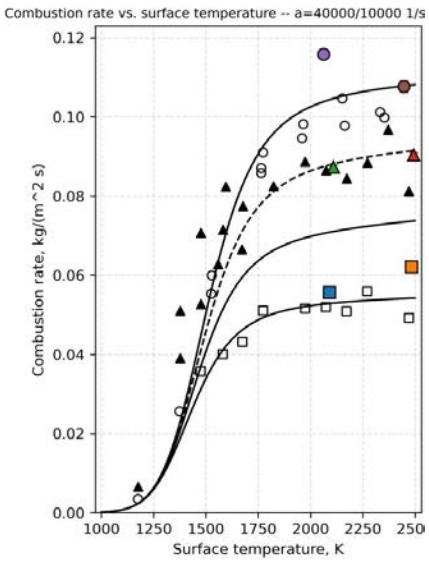

CM³C

Q1-3. Dual Rod – Effect of Radiation

- Two Rods (both 2 mm in diameter)
  - Rod 1: Fixed $T_s = 2100\text{K}$
  - Rod 2: Heat equation solved: $T_0 = 300\text{K}$
  - Surface-to-surface radiation ($\epsilon = 0.8$)
- Gas
  - $U_\infty = 1$ m/s, $T_\infty = 300\text{K}$
  - 20.5% O_2, 78.5% N_2, 1%H_2O by mole
  - Radiation loss using gray emission coefficient fits from RADCAL.
- Solution shown at 30 ms (11,000 time-steps with new method)

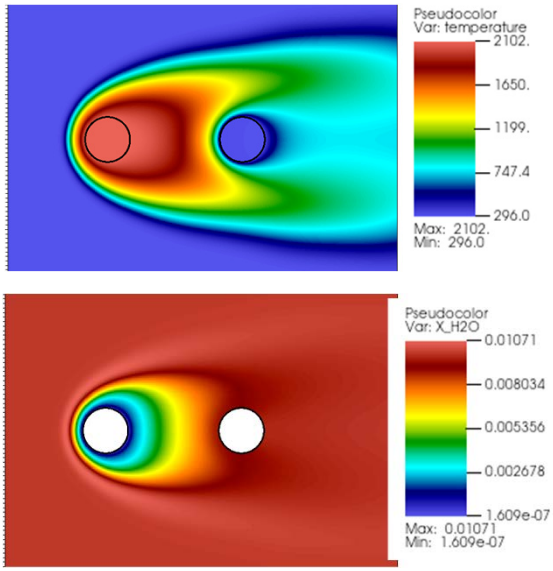

CM³C

Q1-3. Dual Rod – Effect of Radiation

- Radiation Budget at 30 ms:
  - Rod 1: Radiation is ~91% of the heat loss (5.8 kW/m)
  - Rod 2: Radiation is ~34% of the energy into the rod (268 kW/m), conduction is 66% of the energy gain (526 kW/m)
  - Gas Emission: 74 W/m integrated over whole domain. (Optically thin)
- Is absorption important of surface radiation between cylinders important?
  - Ray tracing showed that ~0.2% of emission from Rod 1 is absorbed *en route* to Rod 2.
- Do not need full RTE for microscale problem.
  - S2S radiation with gas emission likely sufficient for microscale sims.

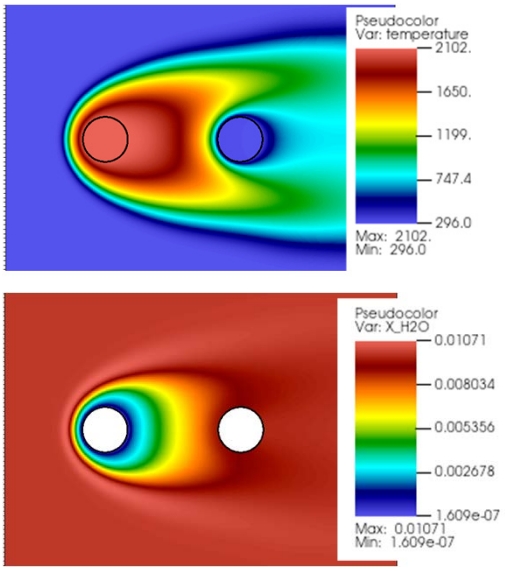

CM³C

CM³C

Questions?



NNSA
National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

CM³C: Center for Multiscale Modeling of Multiphase Combustion

Extreme-scale Many-physics CFD

Spencer H. Bryngelson

The University of Texas at Austin

Oden Institute for Computational Engineering and Sciences

The University of Texas at Austin | UF UNIVERSITY of FLORIDA | PURDUE UNIVERSITY

Every box runs in MFC

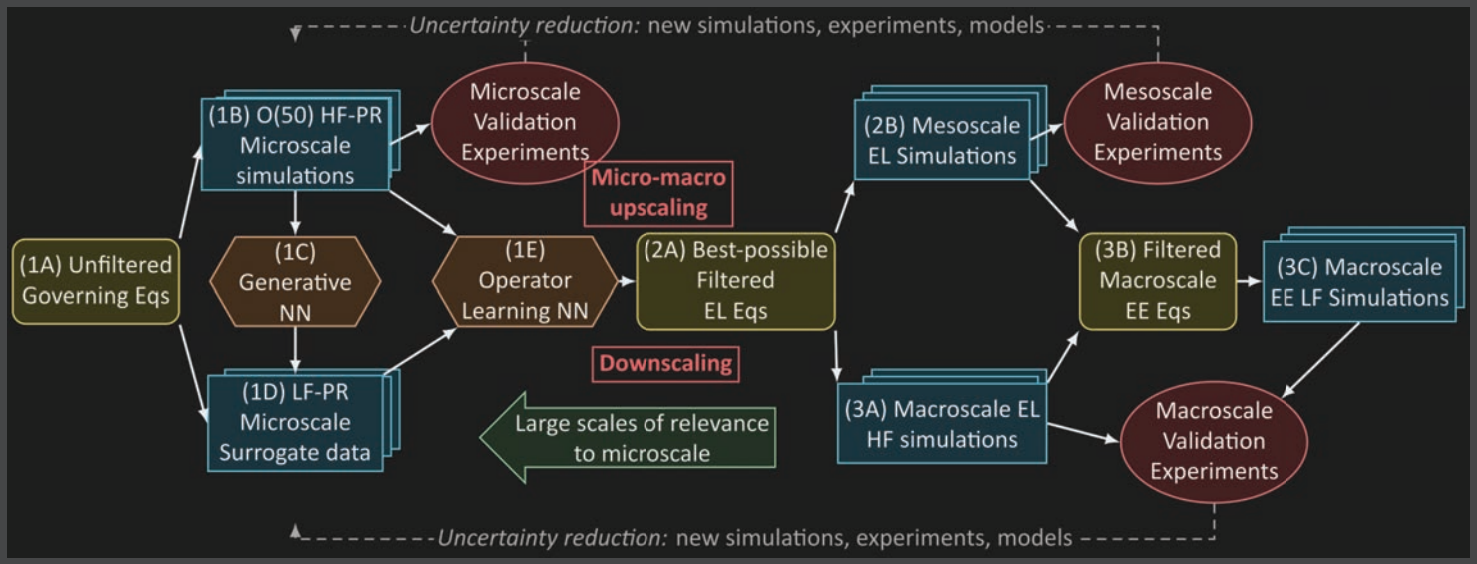

Target: explosive dispersal of burning particles

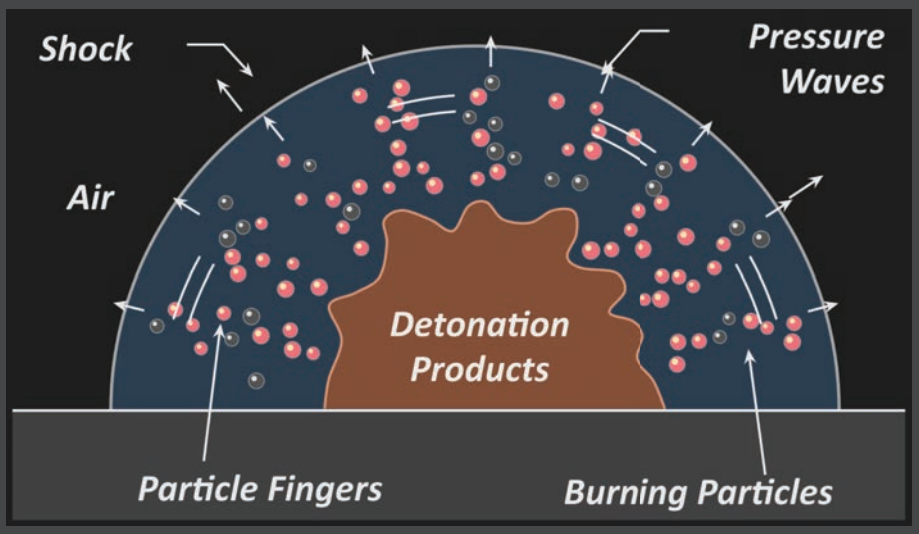

Demonstration problem

HE charge in C/Mg particle shell

10^{11}-10^{13} particles

10^{12} cells, 1 m^3

Quantities of interest

Blast, pressure, volume fraction

Temperature, emission, fingers

Since the April TST

331 PRs merged into MFC
(289 from center team)

11 releases: v5.2.2 to v5.7.1

188K line turnover, 34 contributors

Fixing compiler bugs before they hit us

Root caused 45 compiler bugs

18 merged LLVM PRs

2 new upstream AFAR releases
(AMD compilers)

Machines

AMD Internal (MI350X +)

Tioga (MI250X)

Tuolumne (MI300A)

Frontier (MI250X)

UFlorida (B200)

Gatech (GH200)

github.com/MFlowCode/MFC, Bryngelson et al. HPCTESTS SC'26 (2026); Wilfong et al., Comput. Phys. Commun. (2026)

New in MFC since April

Particles

- Collisions
- Second-order IBM
- Reacting surfaces
- Per-species ghost cells
- Moving bodies (cut-volume)
- Sub-cell wall conduction
- Lumped particle temp.
- Forces from face fluxes

Reacting flow

- QSS stiff chemistry
- Heterogeneous surface reactions
- Reactive burn/burning surfaces
- Mie-Gruneisen, JWL, Vinet EOS
- Wall treatment (e.g., conduction)

Euler–Lagrange

- Drag, added mass, two-way coupling

Scale / Speed

- AMR: multilevel, subcycled, etc.
- Low-Mach solve
- Locality-aware IBM
- Streaming CFD

github.com/MFlowCode/MFC

MFC at scale: weak scaling

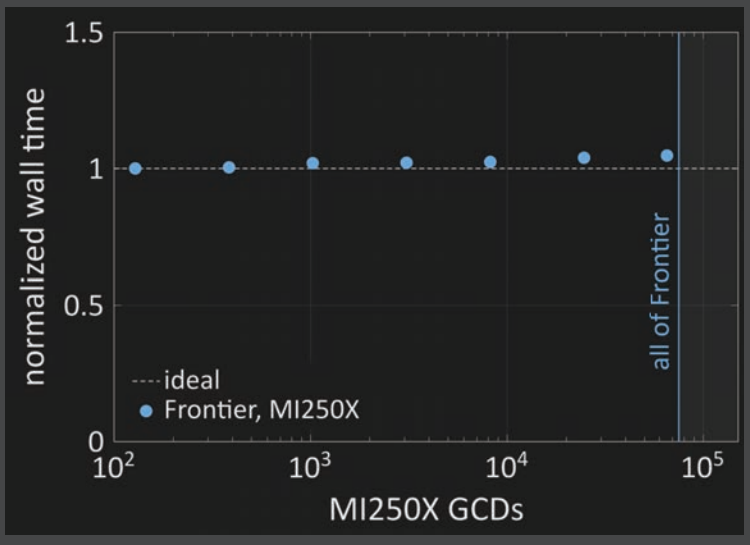

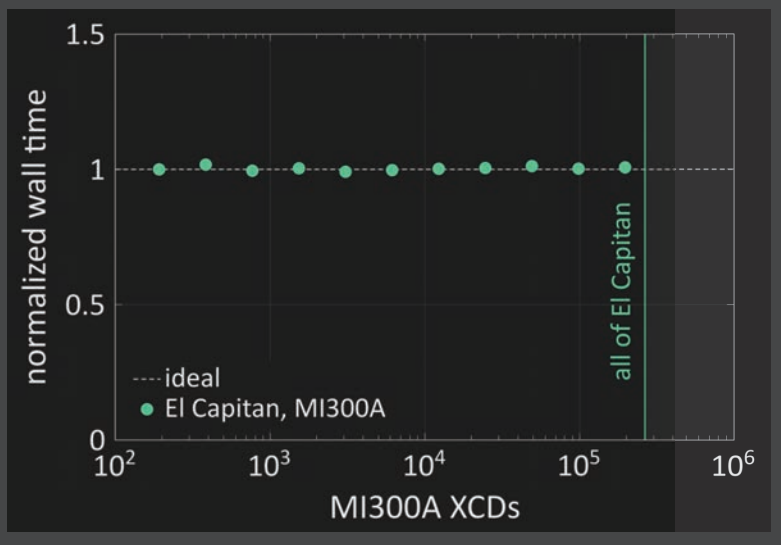

Near-ideal weak scaling; Leveraged DOE INCITE + ALCC + NSF ACCESS-CI allocations

Wilfong et al., Comput. Phys. Commun. 322, 110055 (2026); Radhakrishnan et al., Comput. Phys. Commun. 302, 109238 (2024)

Roadmap: Year 1 targets

Item	Proposal, Y1	Now
Equation of state	Ideal gas	Ideal, Mie-Gruneisen, JWL, Vinet (Y2 item)
Reactions	O(1)	Any mechanism; surface/condensed-phase, stiff time stepping
Particle contact	None	Soft-sphere collisions (Y2 item)
Load balancing	“Simple"	Locality-aware IBM; AMR+Load balancing/subcycling (Y3+ item)
Grid	10^{12}	62T grid cells, 256 Tuo nodes [streaming], 86T cells [1K nodes]
Resolved particles	n/a	1.2B on MI250X Frontier
EL particles	10^{10}	EL particles across ranks; ~10^{11} on Frontier

Proposal Table 2.

Particle-resolved reacting flow

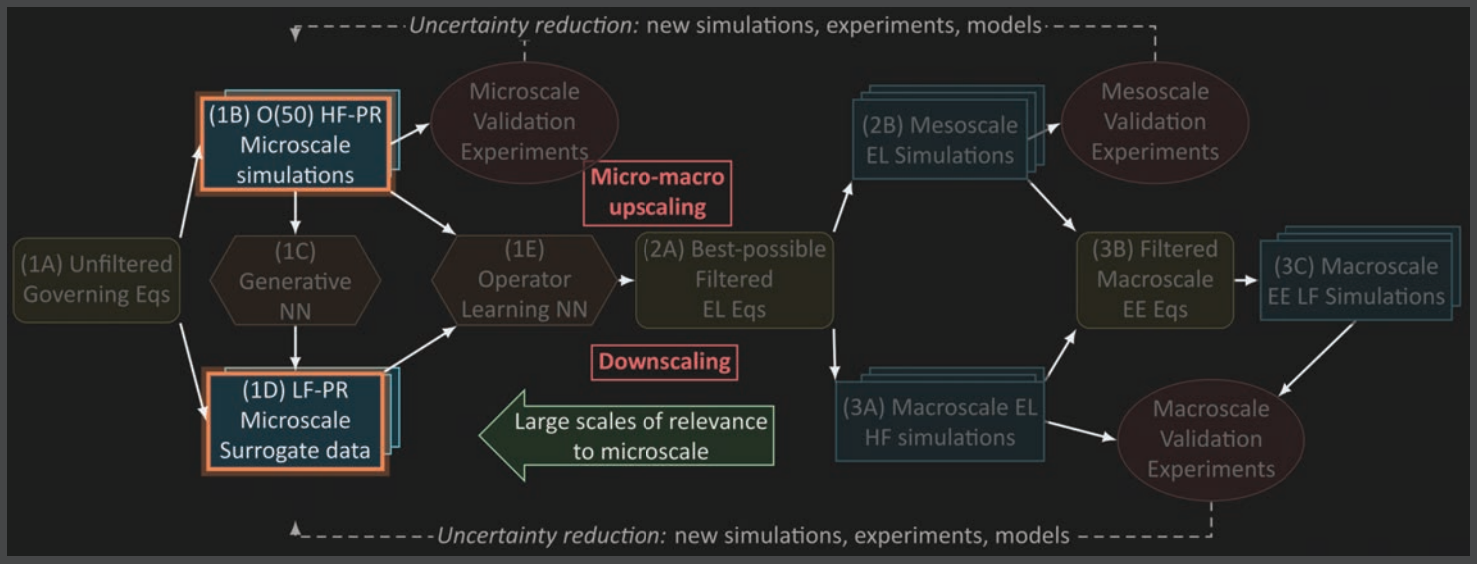

Immersed boundary treatment

1.2 billion resolved particles

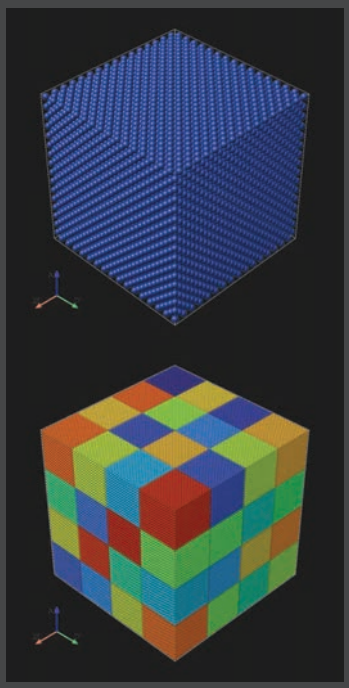

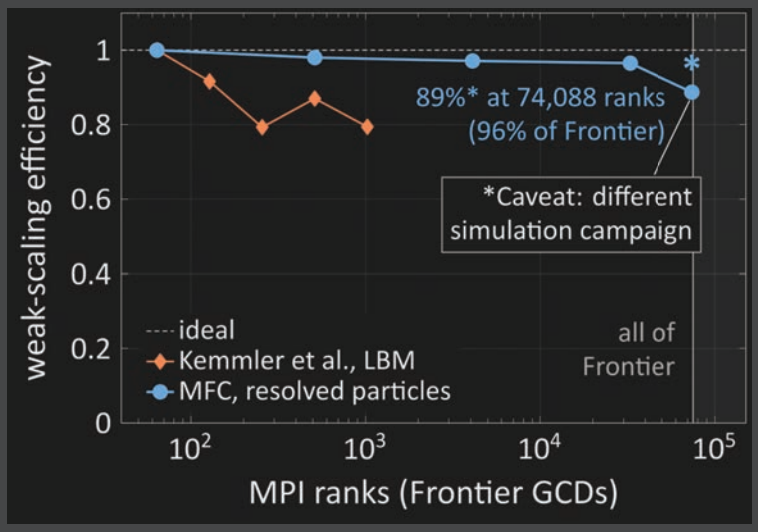

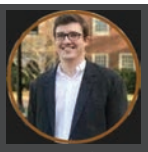

Dan Vickers
Georgia Tech

16,384 particles/rank x 74,088 ranks

= 1.2B IBM resolved particles

LBM reference: Kemmler et al. IJHPCA 2025

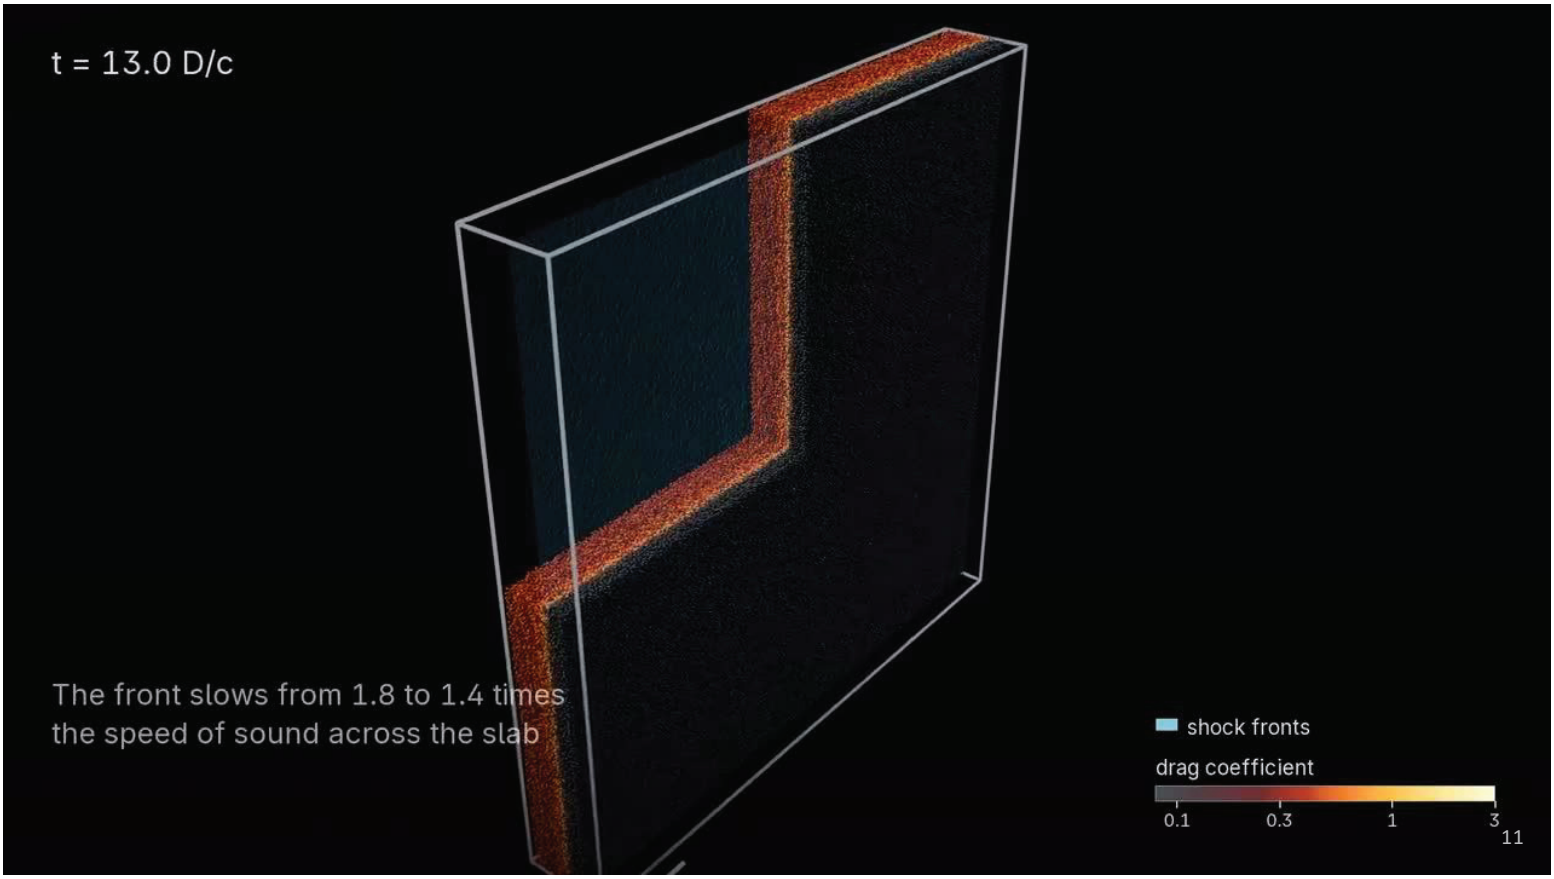

Particle-laden shell

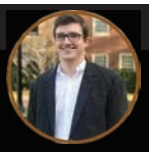

Dan Vickers
Georgia Tech



Ben Wilfong
Georgia Tech

Vickers, Wilfong, Bryngelson; APS Gallery of Fluid Motion (2026)

Into reacting flows…

Immersed boundaries for reacting particles

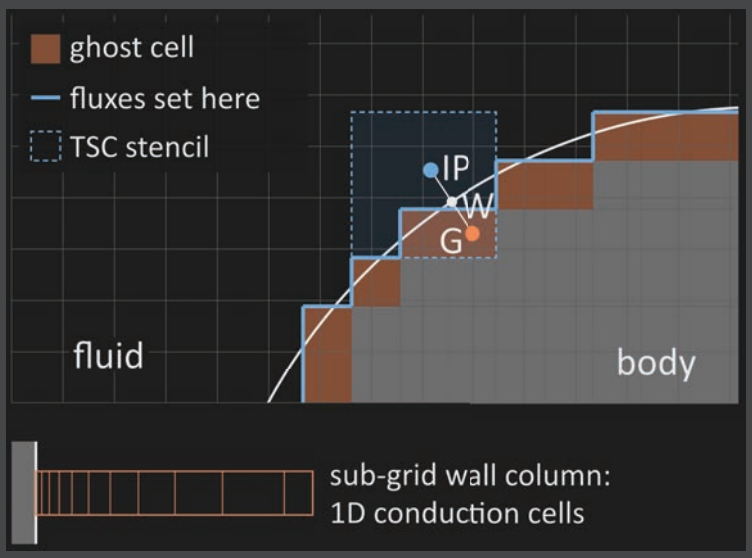

Beyond ‘traditional’ IBM

Second-order accuracy for all qtys/fluxes

- TSC image points: second-order wall T / heat

Conservative reacting walls

- Zero-mass-flux wall: gas gains what gasifies
- Species conserved to round-off
- Per-species ghost weights

Moving bodies

- Cut-volume + GCL + moving flux wall

Sub-cell thermal layers

- Wall columns: within 2.1% of reference

Forces from face fluxes

Reacting-flow verification: 3D

Problem	Dim.	MFC result	Compared with
Flame hit by vortex ring	3D	**2nd order**	H2O mass, flame area
Flat H2-air flame	3D	**0.05%**	flame speed, Cantera
Burning carbon sphere, CO flame	3D	**0.3%**	burn rate, 1D spherical reference
64 burning particles	3D	**+2%**	film theory, grid converged
Element conservation [burning array]	3D	**10^{-12}-10^{-9}**	O, H, N in closed box

13+ more in 0D-3D: backup

Runs: Tuolumne, Tioga, Frontier. Chemistry: Cisneros-Garibay, Bryngelson et al., Comput. Phys. Commun. 320, 109987 (2026)

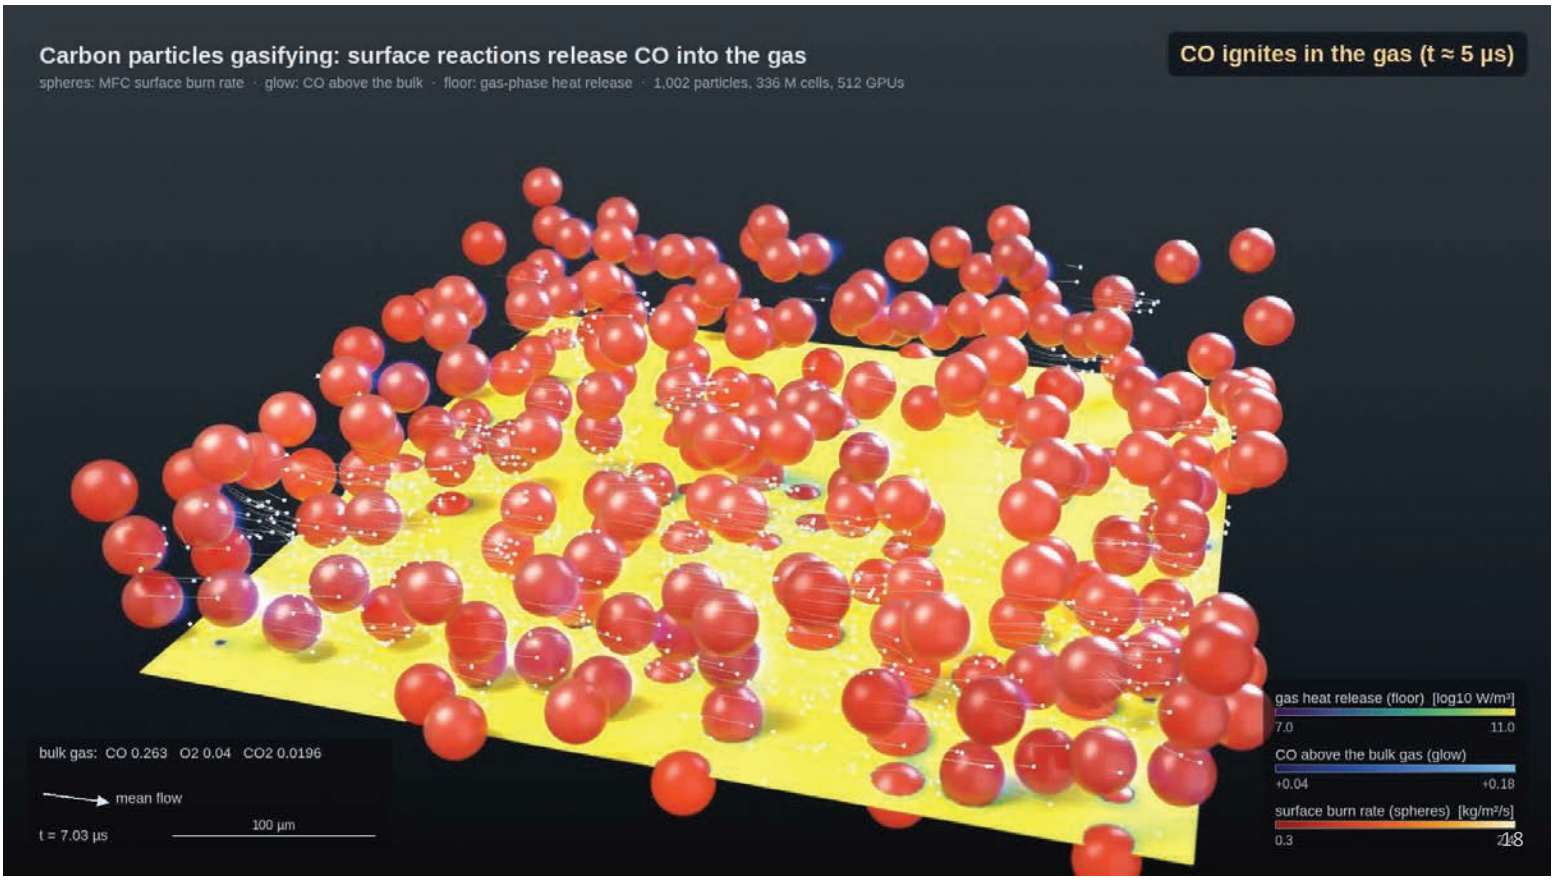

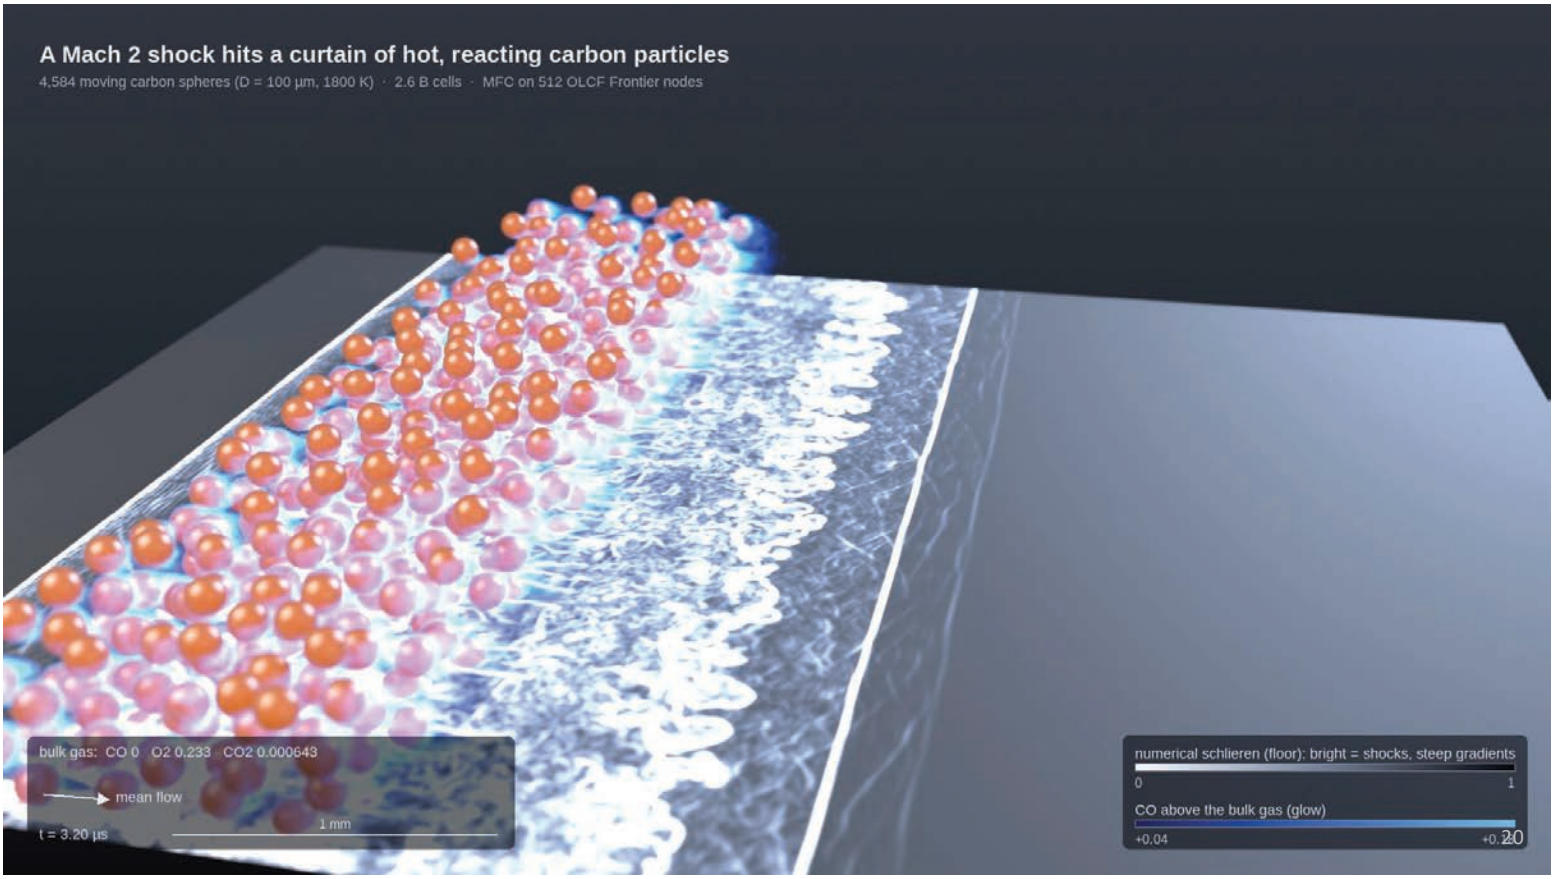

Particle burn rate

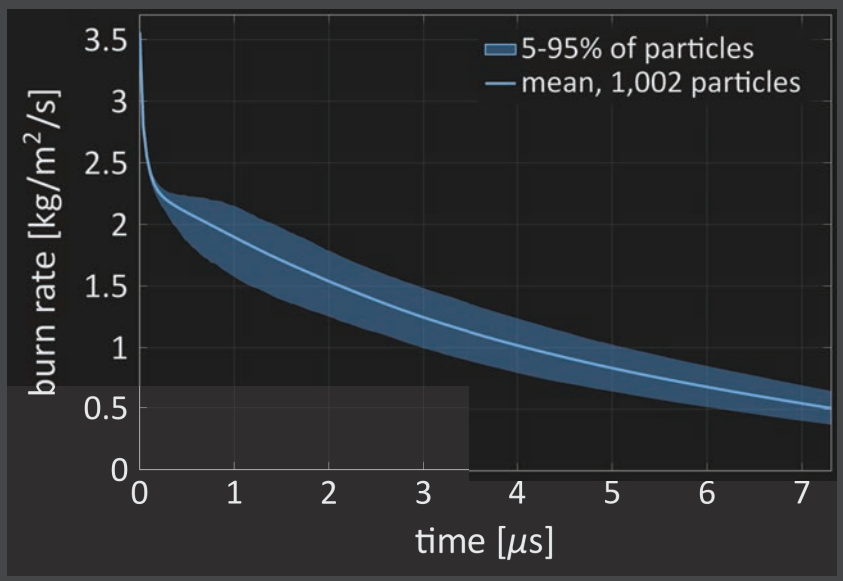

1,002 particles tracked

3.6 kg/m^2/s at start; 0.5 at 7.3 us

Falls as O2 consumed

Reference

Within 2% of film theory

Input to EL closures

Burn rate vs local gas state

Burning carbon shell dispersed by a blast (2D)

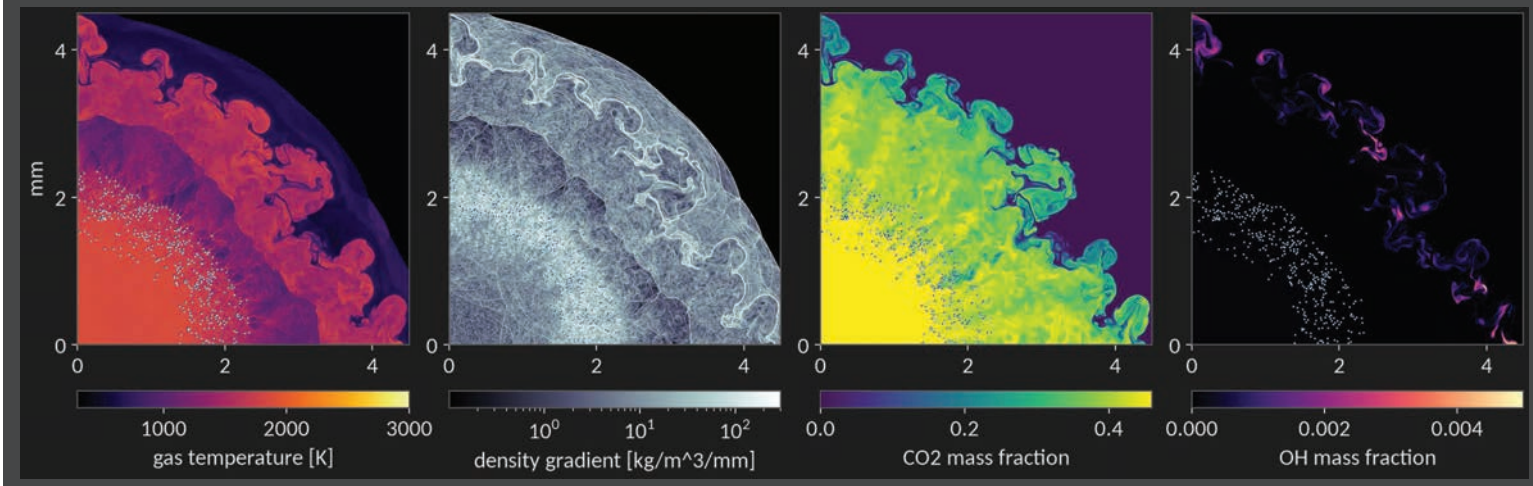

Blast wave; shell breaks into jets

Particles accelerate to 540 m/s

Products afterburn with air (OH)

Hot carbon surfaces react

2D section, 392 particles (25 um), 12 cells/D; 250 bar charge.
More detail: D. Vickers [GT/UT], Day 2

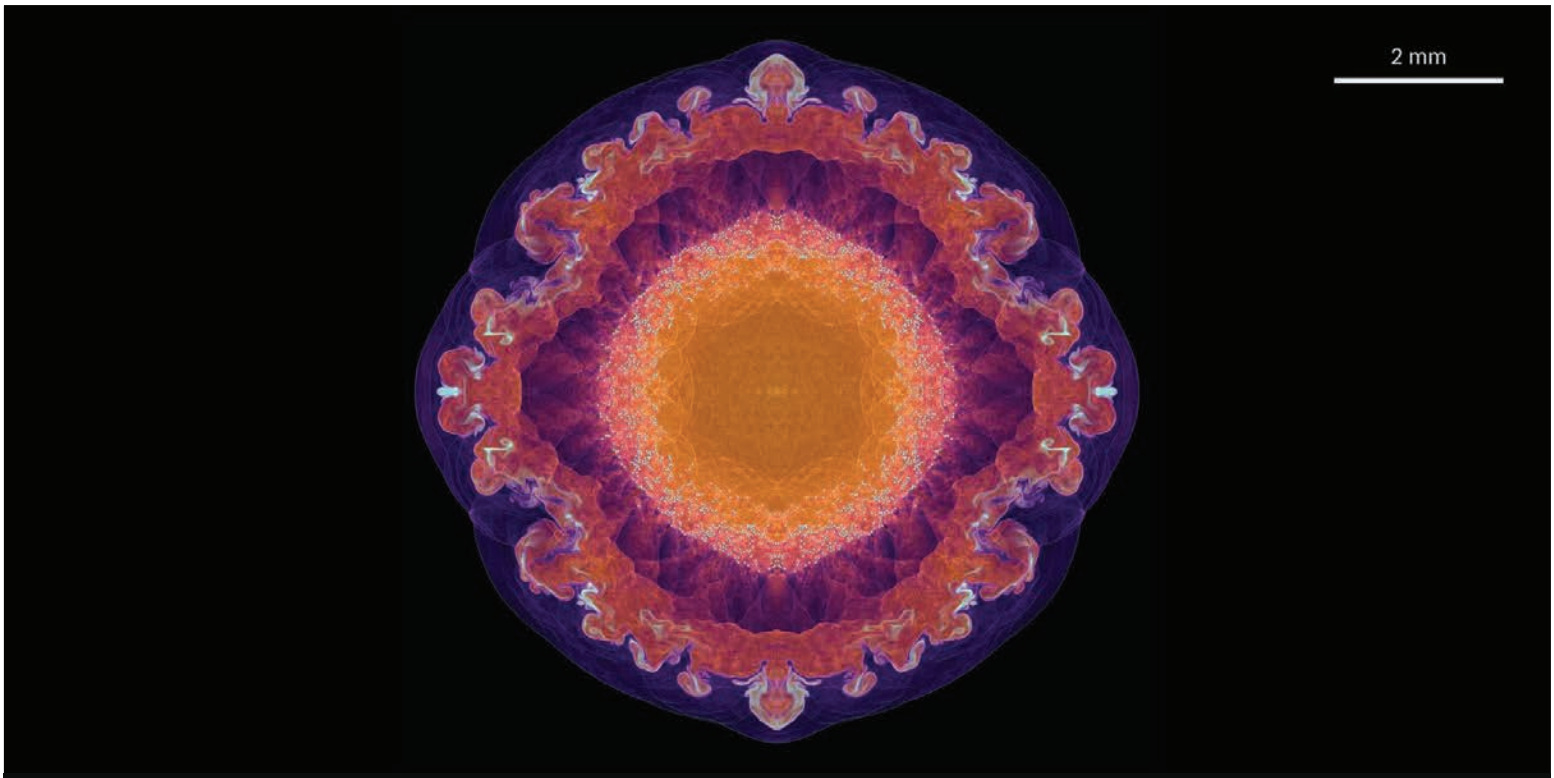

Shell, t = 3.5 us. Gas temperature with Schlieren; cyan: OH; dots: resolved carbon particles

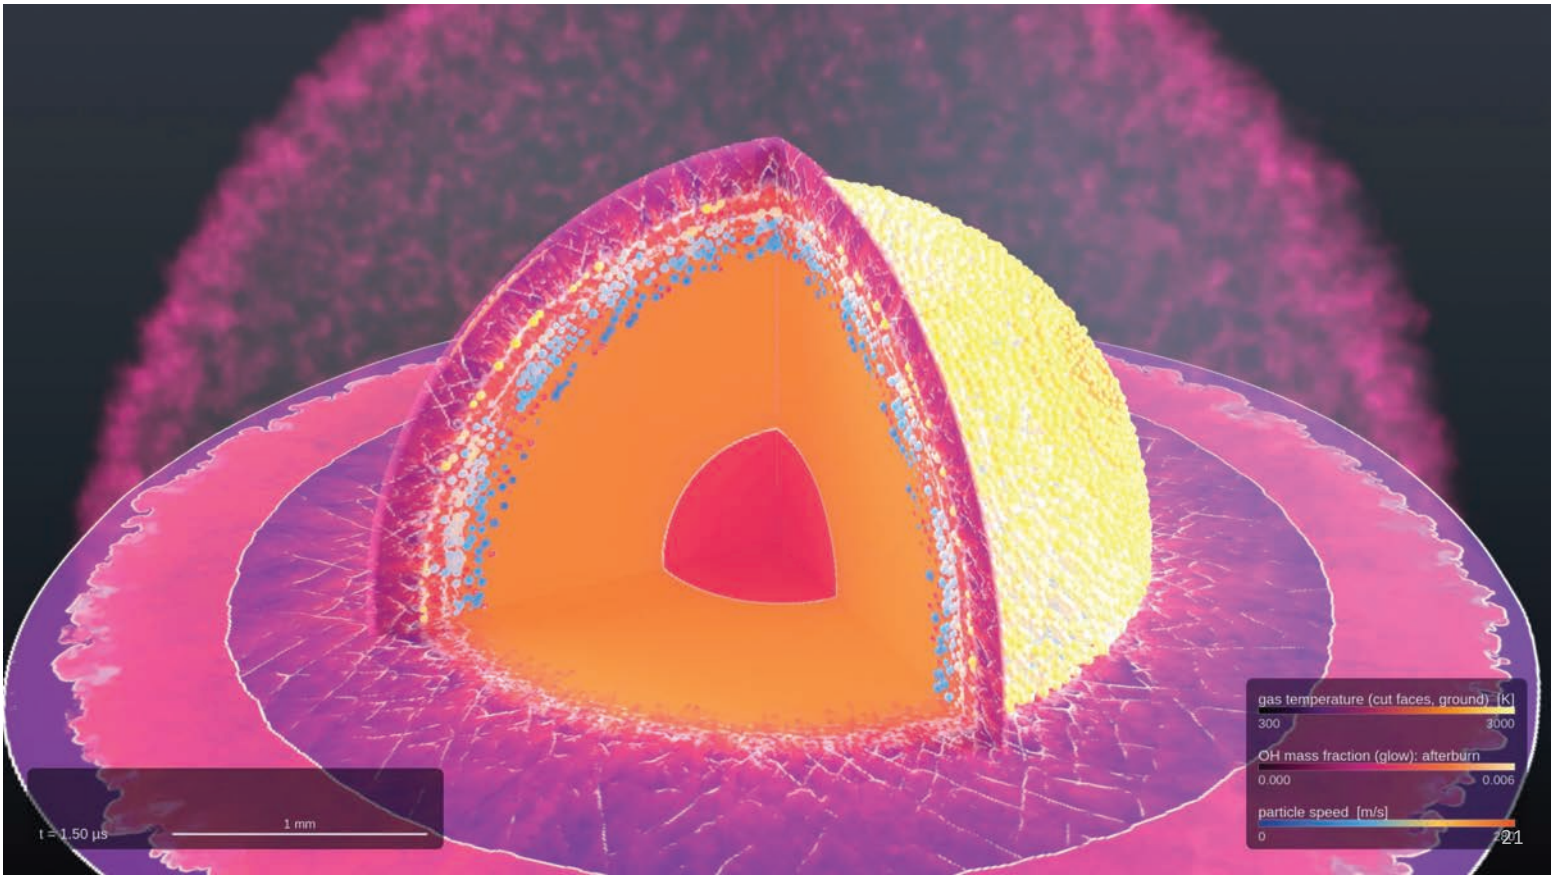

21

Same computers, larger+faster simulation

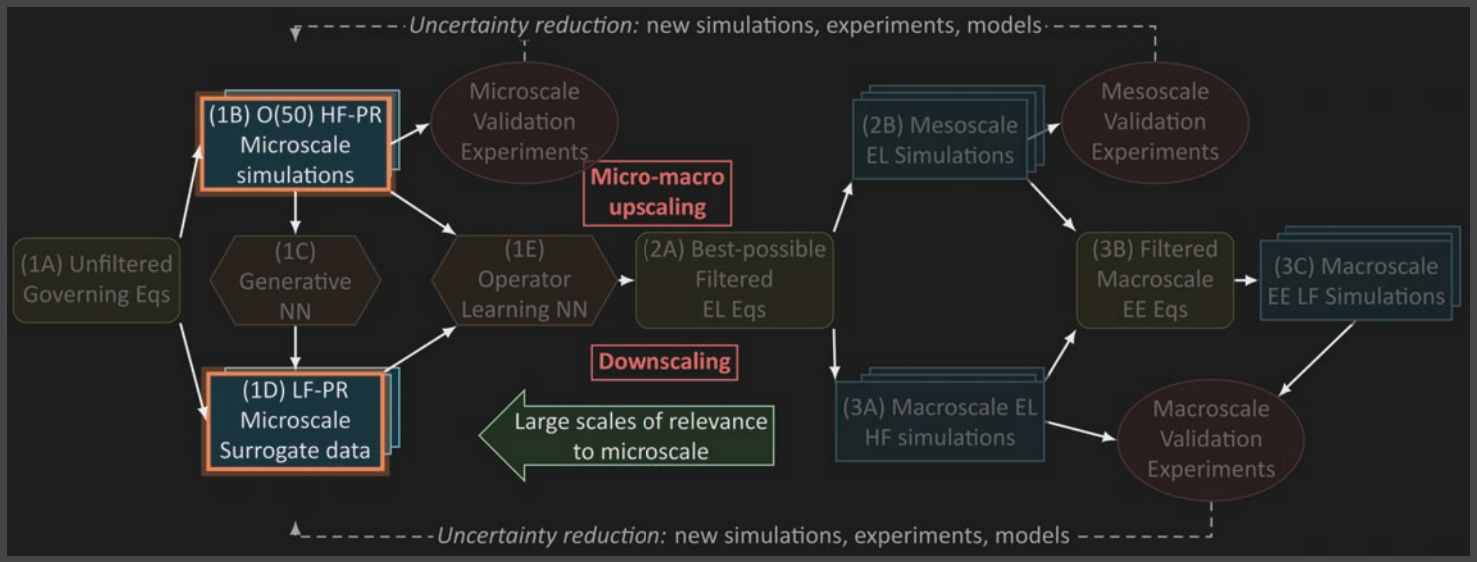

9

Low Mach treatment

Low-Mach projection solver

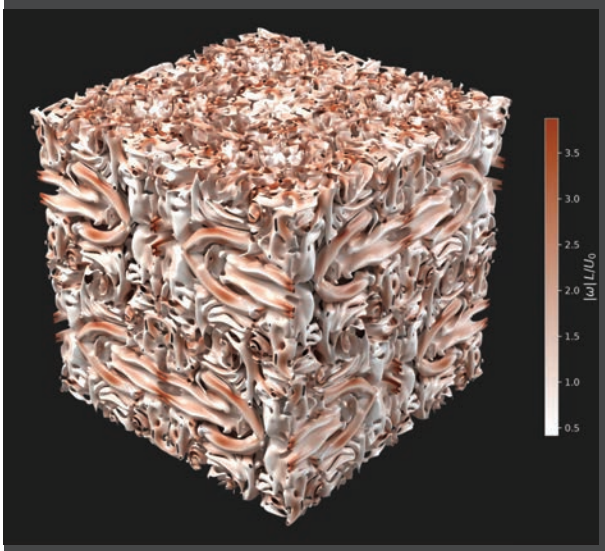

Shown: Vorticity magnitude

Explicit transport, implicit pressure

Multigrid-preconditioned CG

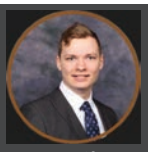

Ben Wilfong
Georgia Tech

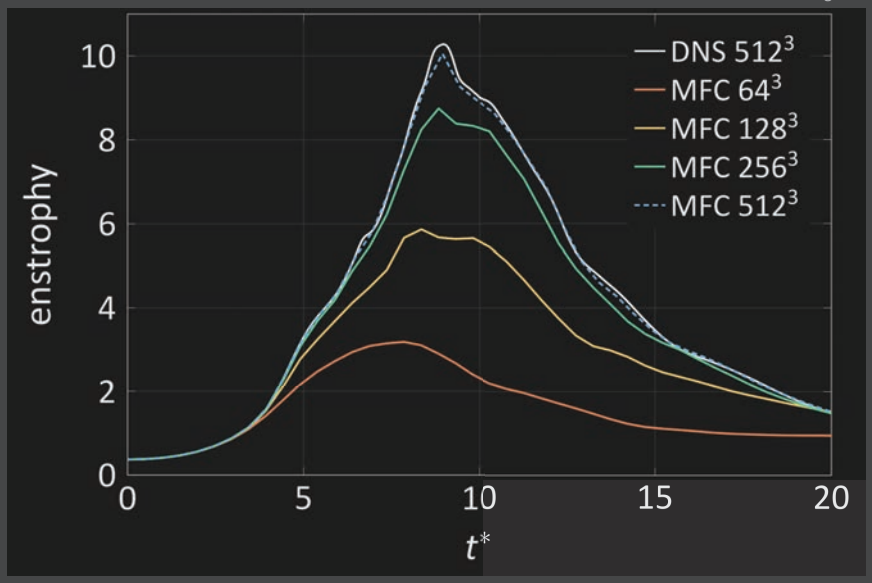

Mach 0.01, Re 1600 Taylor-Green vortex

Low-Mach with chemistry: burning carbon rods



Ben Wilfong
Georgia Tech

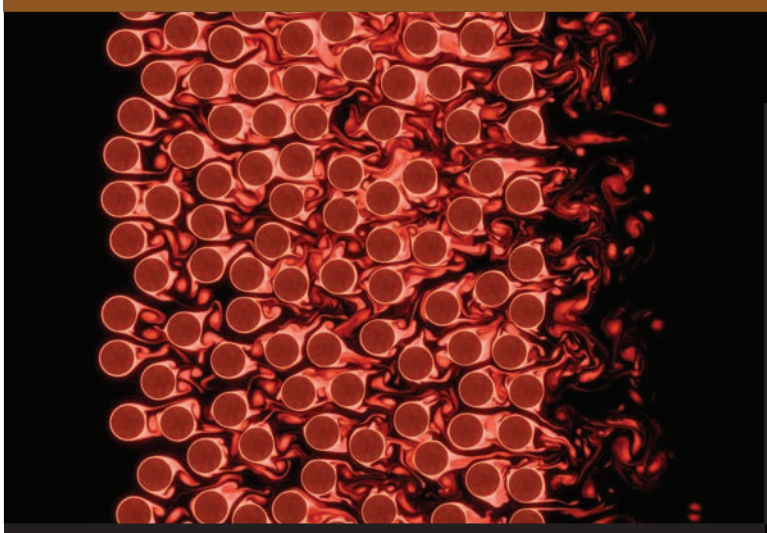

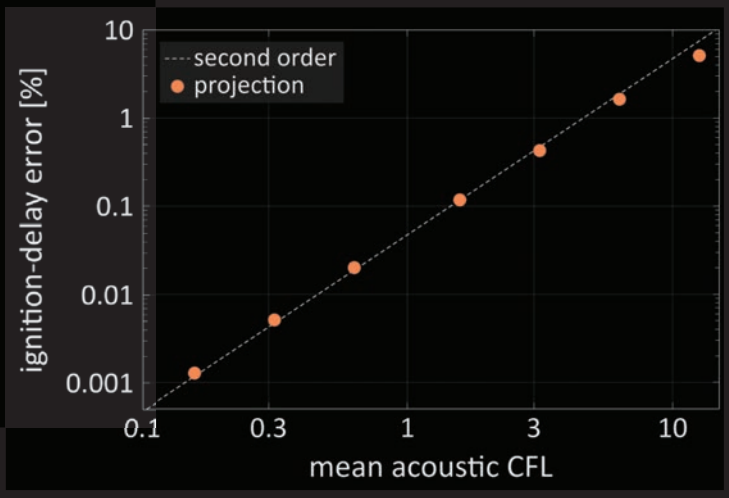

Mach 0.05, Re 2,400

90 rods, 192 cells/D

11 species, 33 reactions

Carbon surface chemistry

Second order accurate

Explicit solve within 0.05% of Cantera

Low-Mach solver at scale



Ben Wilfong
Georgia Tech

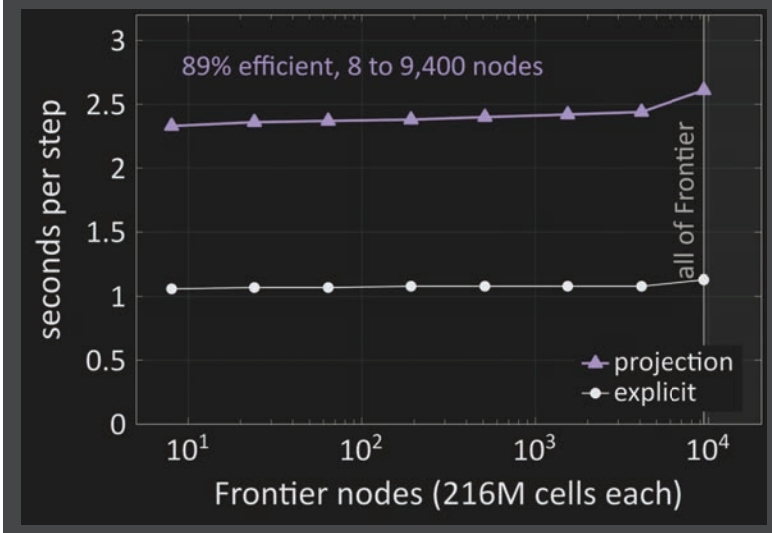

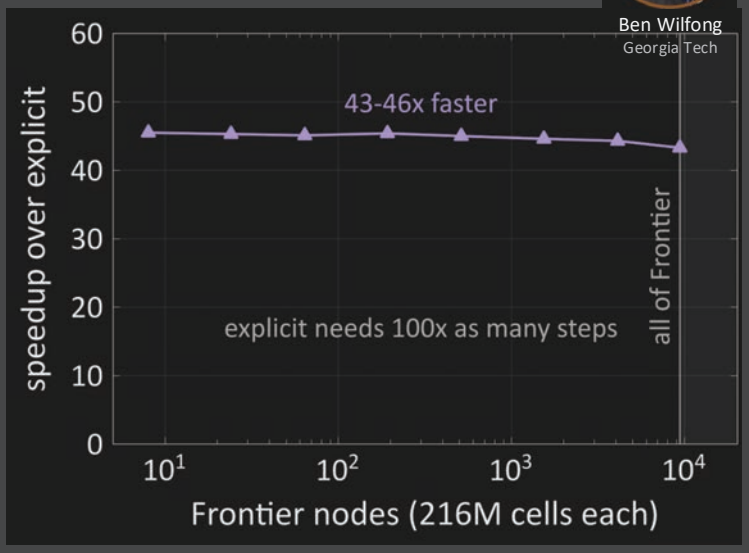

Near ideal weak-scaling, 45x faster than explicit @ Mach 0.01

Above: 216M cells per node, explicit is 100x more steps

Streaming CFD

Streaming from flash: How it works

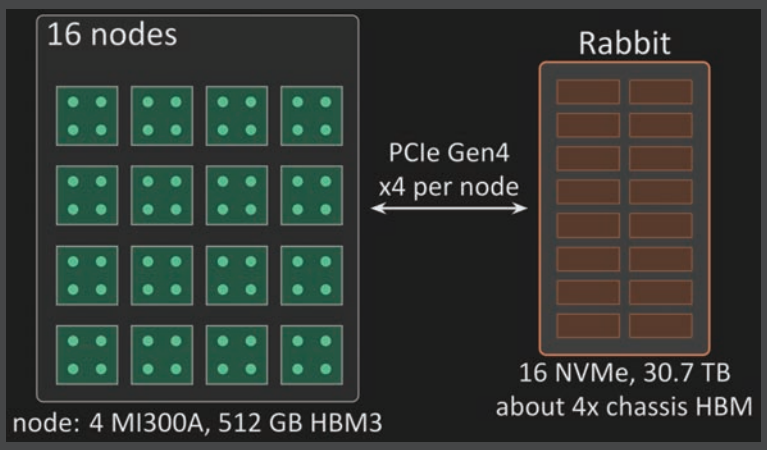

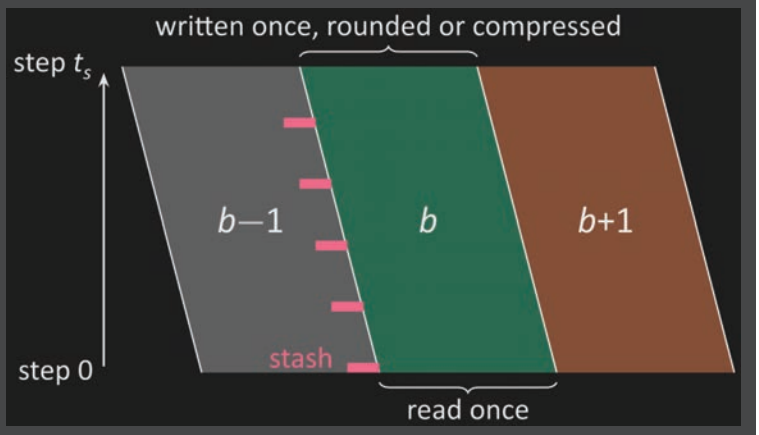

Whole state on the Rabbit

Blocks stream through HBM

Streaming: read, advance, write at once

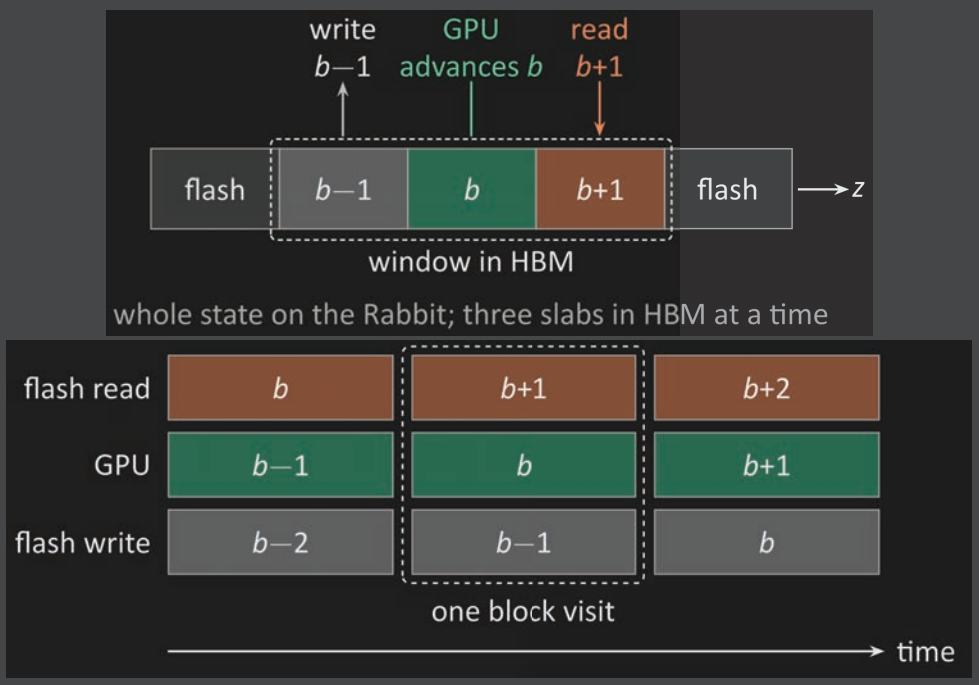

Flash I/O hidden behind GPU work

Beyond GPU memory: streaming from flash

Record CFD sim: 200T pts on Frontier

1 node: 146B cells, 96% of in-memory (regular MI300a) speed

Compressed: 670B cells/node, 94% of 'regular' speed. Custom hip algo. added to `zfp`

256 Tuo/El Cap nodes: 62T cells 100% El Capitan: 30 quadrillion Speed limited by compute!

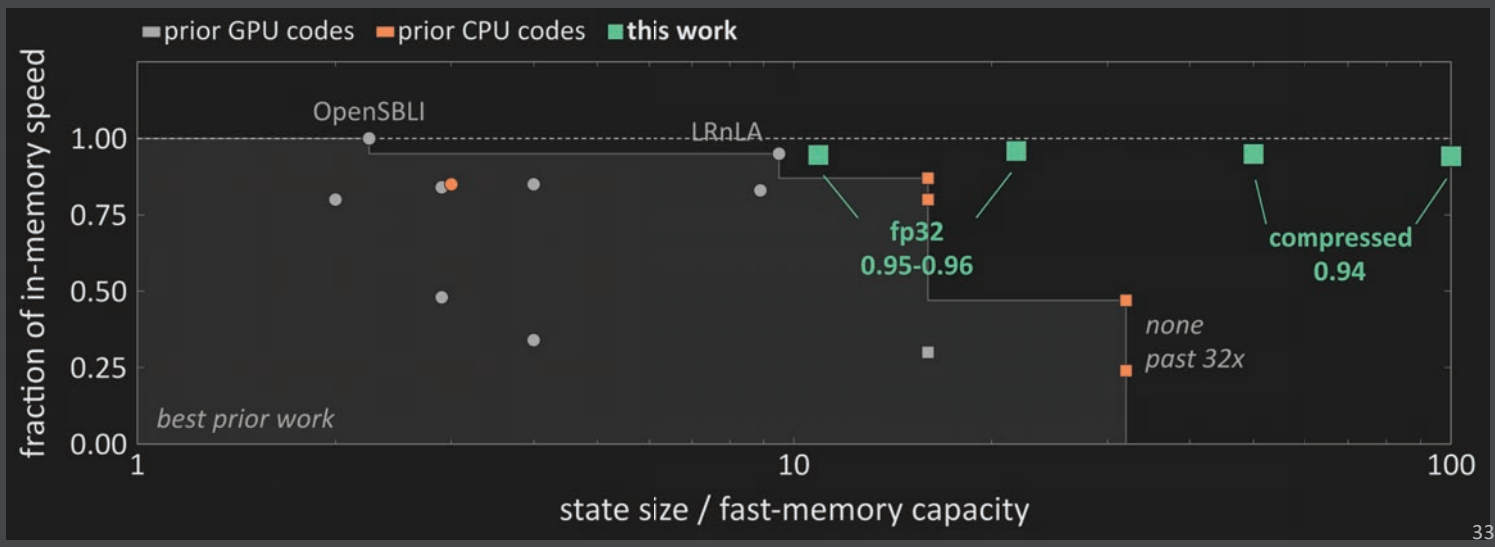

AMR + Load Balancing

AMR: Mach 2.5 shock through moving 2-way coupled spheres

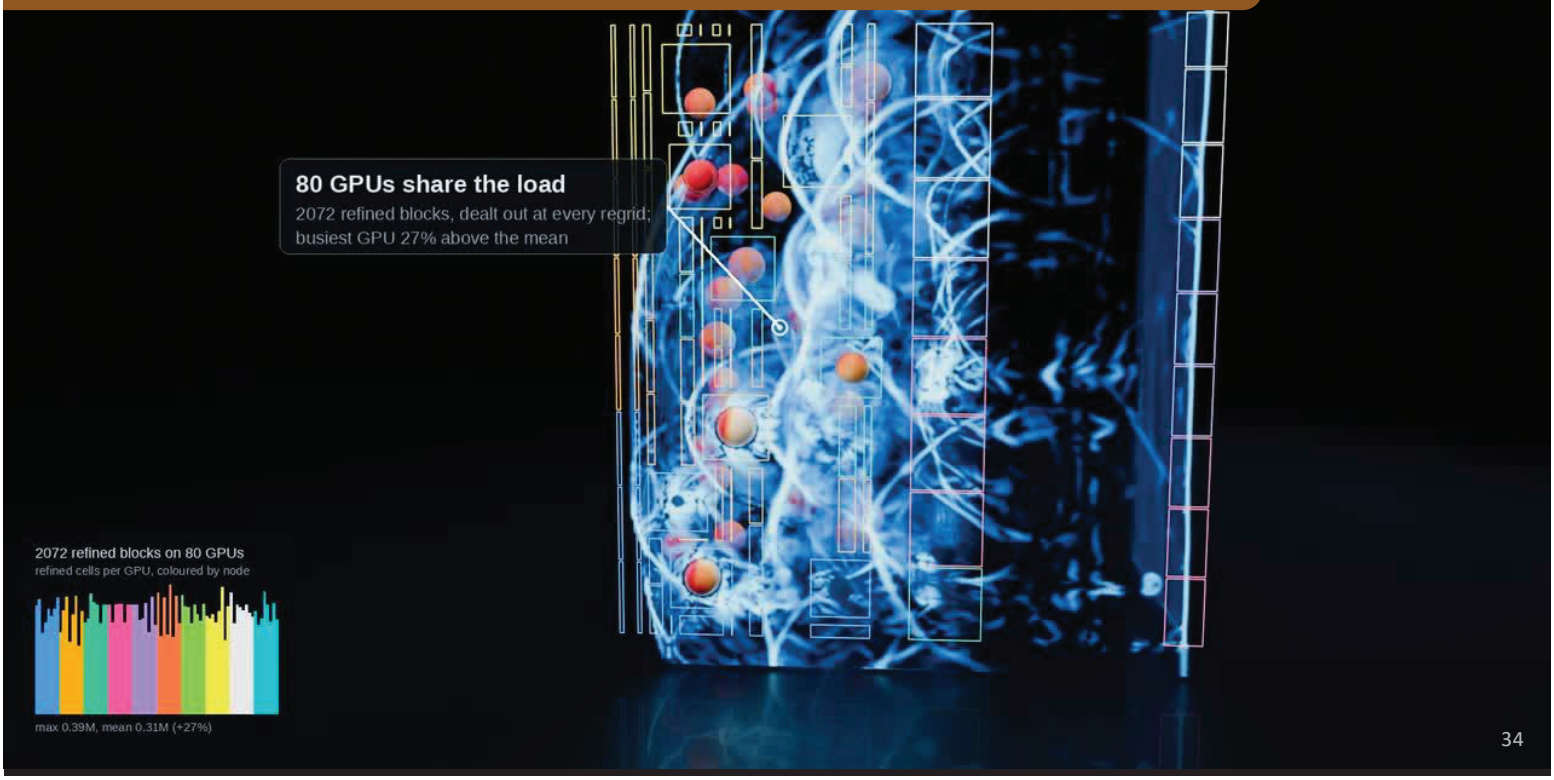

AMR Performance: MFC vs PeleC

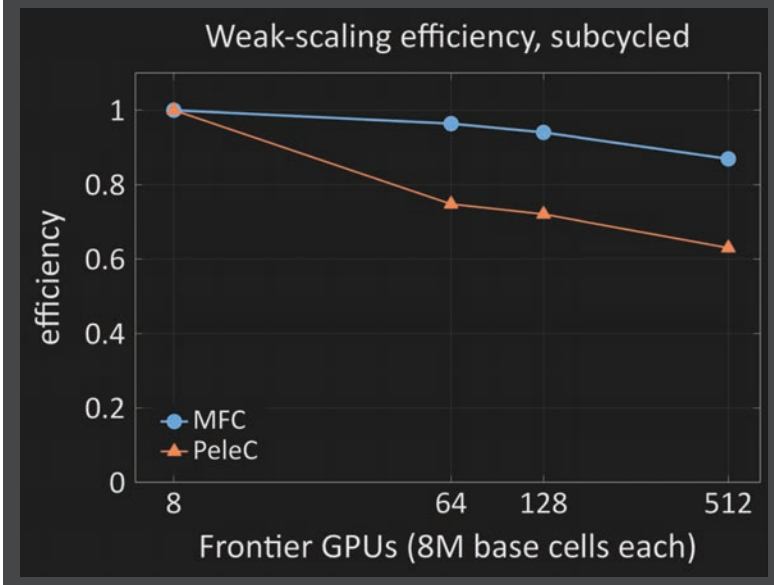

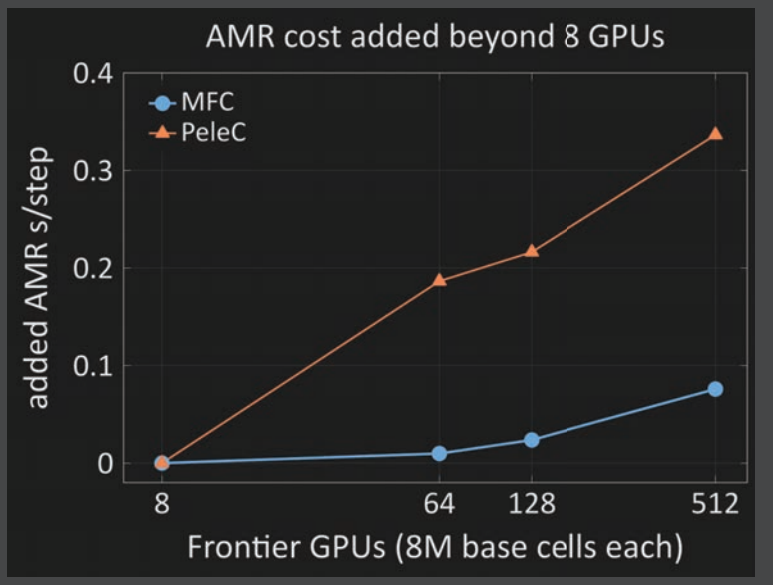

Subcycled efficiency @ 512 devices: MFC 0.87, PeleC 0.63

Case: domain filling blasts, 8M base cells/GPU

Subgrid particles

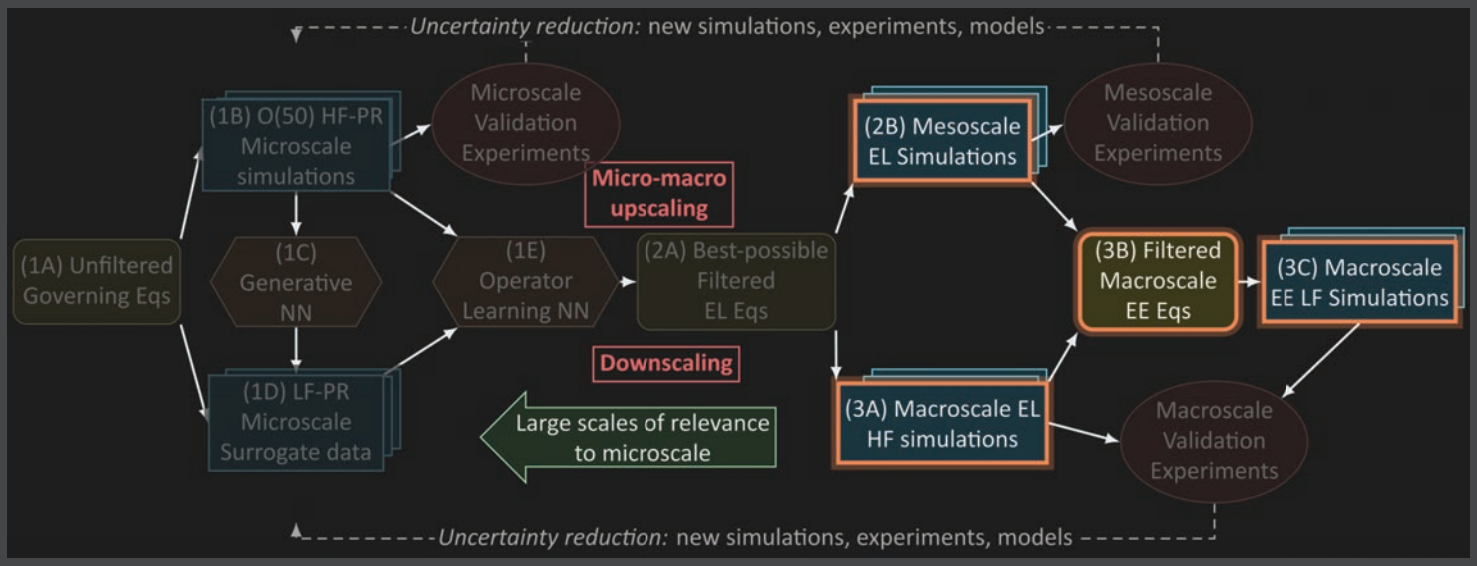

Euler-Lagrange: Mach 10 blast, 500K particles

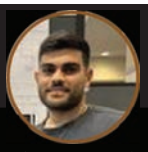
Thierry Daoud
U. Florida


Jose Aguilar
Alum. U. Florida


Ben Wilfong
Georgia Tech

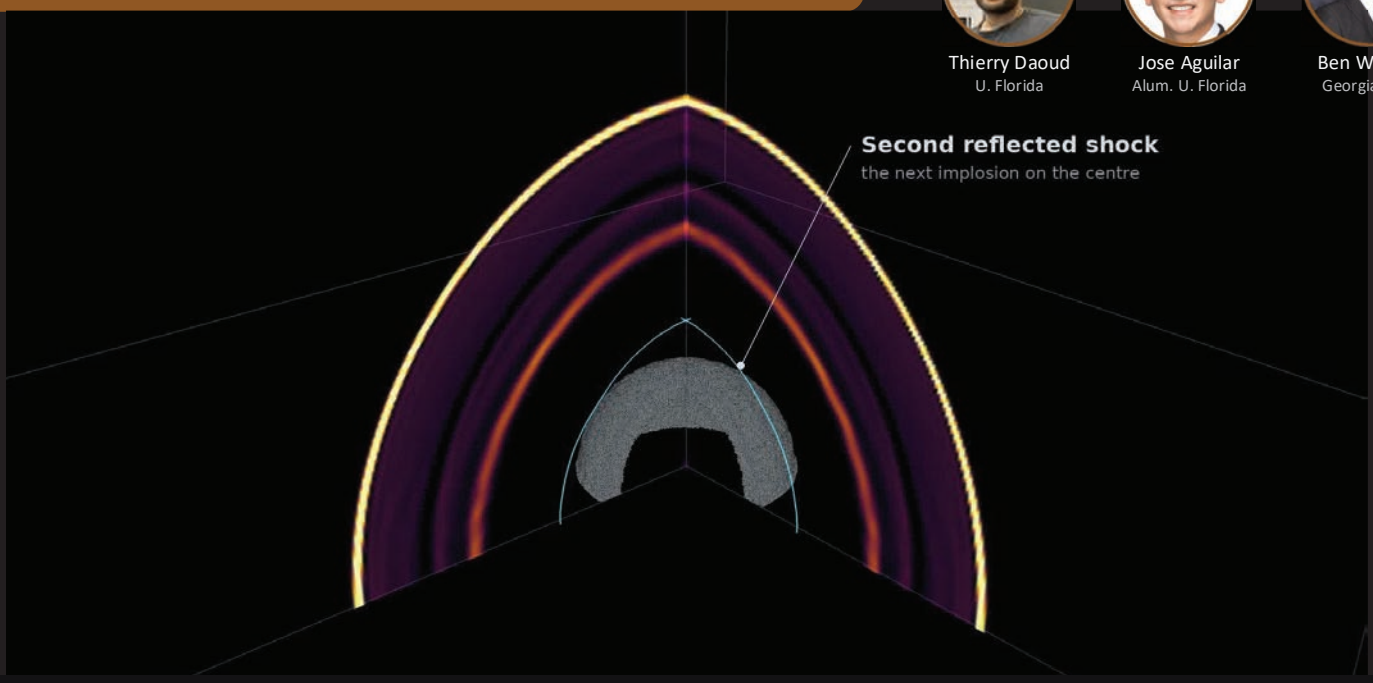

Capacity: 1M particles/GPU (Tuolumne); x 75,264 GCDs = ~10^{11} on Frontier

Fourth-order kinetic moment closures


Rodney Fox
Iowa State

Particle velocity distribution f(t, x, u)

$$\partial_t f + u \cdot \nabla_x f = C_{BGK}(f)$$

Free transport + BGK collisions (Knudsen number Kn)

Moments to 4th order: 35 (3D velocity)

$$m_{ijk} = \int u_1^i u_2^j u_3^k f \, du \qquad \partial_t M + \nabla \cdot F(M, M_5) = S$$

HyQMOM closure for the 5th-order flux

$M_5 = M_5(M)$ from orthogonal polynomials of f

Hyperbolic; realizable $\Delta_k \geq 0$

Two particle clouds head-on collision at Mach 5, 3D

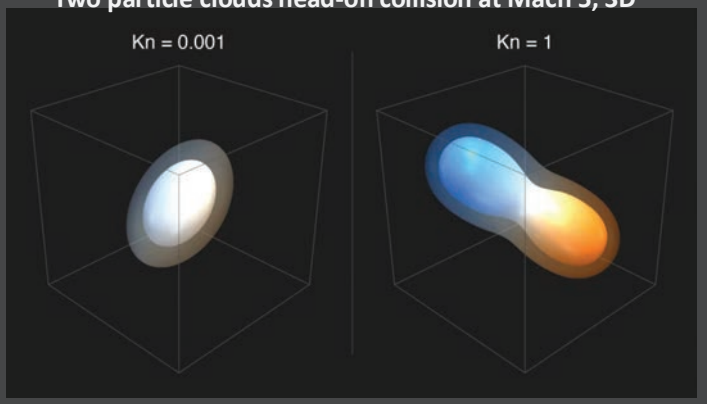

Low Kn: collide & merge. High Kn: pass through!

Bryngelson, Fox, Laurent, J. Comp. Phys. (2026)

Euler–Euler particle clouds possible w/ anisotropic, non-equilibrium velocity distributions

ML closures

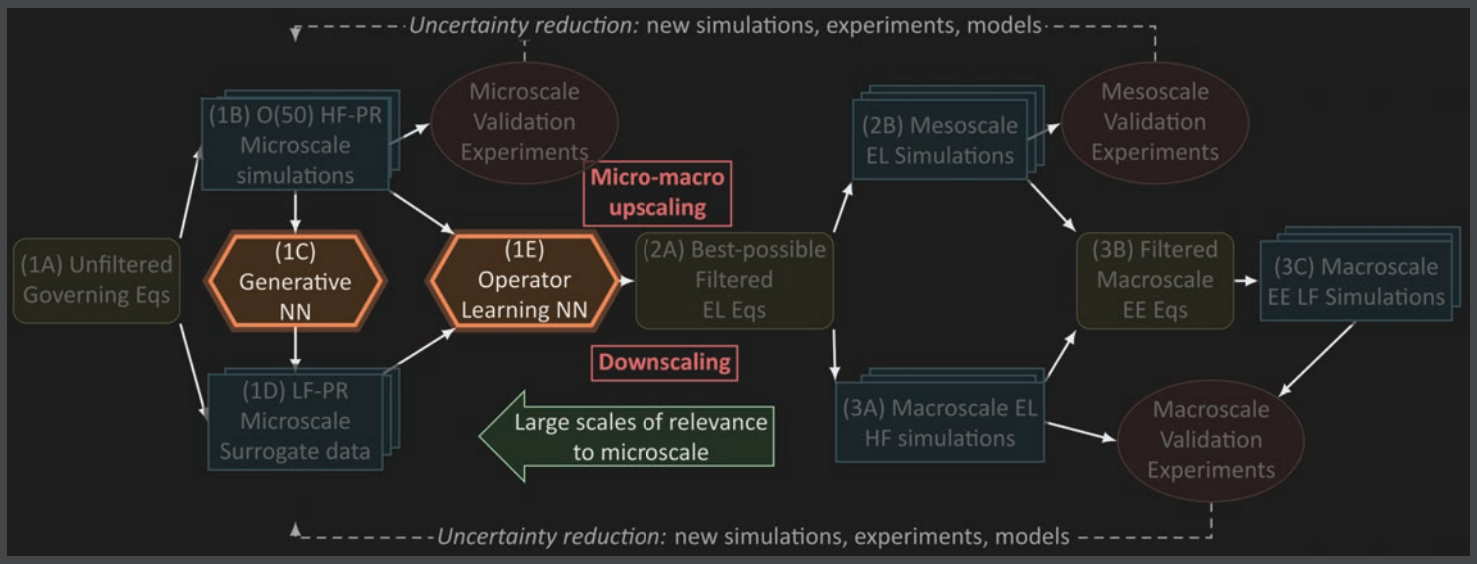

roseNNa: trained networks inside MFC

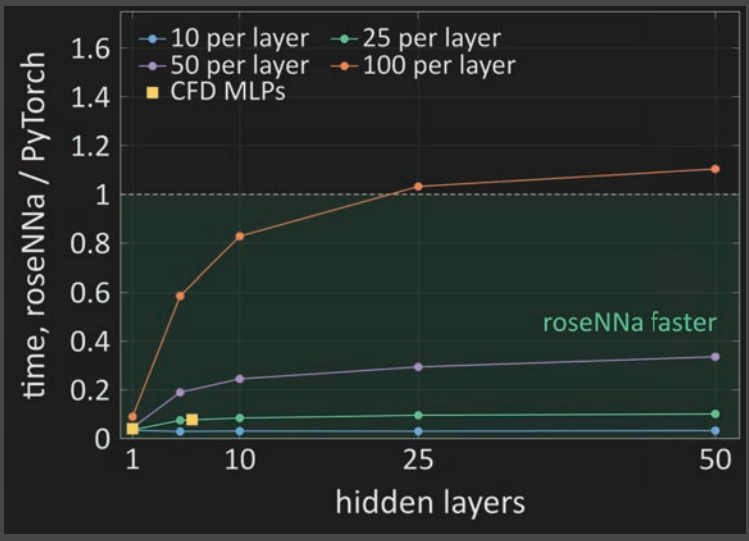

v1.0 (Sept 2026)

ONNX in; Fortran 2008 or C11 out

No allocation; pure, callable in GPU loops

MLP, CNN, LSTM, GRU

Validated:

NVIDIA (NVHPC), AMD GPU (amdflang)

Faster than PyTorch

… for shallow CFD networks

github.com/comp-physics/roseNNa | Bati + Bryngelson, Comput. Phys. Commun. 296, 109052 (2024)

From simulation to closure

Data out (MFC)

Resolved curtains, burning arrays

Per particle, save: force, heat flux, burn

Local gas state of particles (EL interpolation, 2nd/4th order)

Tables: particle x time x (features, targets)

Train (Harley, Zare)

Closure for drag, heat transfer, burn rate, etc.

ML framework trains

Export model to ONNX

Model in (MFC + roseNNa)

ONNX to Fortran/C, no allocation

Called inside GPU kernels;

Drops into EL drag, heat, and burn-rate calls (etc.)

MFC side ready: per-particle outputs written; roseNNa makes trained closures GPU-callable in the EL source terms

github.com/MFlowCode/MFC and github.com/comp-physics/roseNNa

Year 2

Demonstration problem

Explosive dispersal of full burning carbon shell

JWL products, collisions, two-way coupling

Boundary layers, group combustion, better models!!!

Physics

Magnesium particles: vapor-phase flames

EL collisions; ML closures in EL/EE runs

Scale

AMR+load-balance [+streaming? maybe?]

CM³C

The University of Texas at Austin

Questions

The University of Texas at Austin

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

Backup

Conservative reacting walls

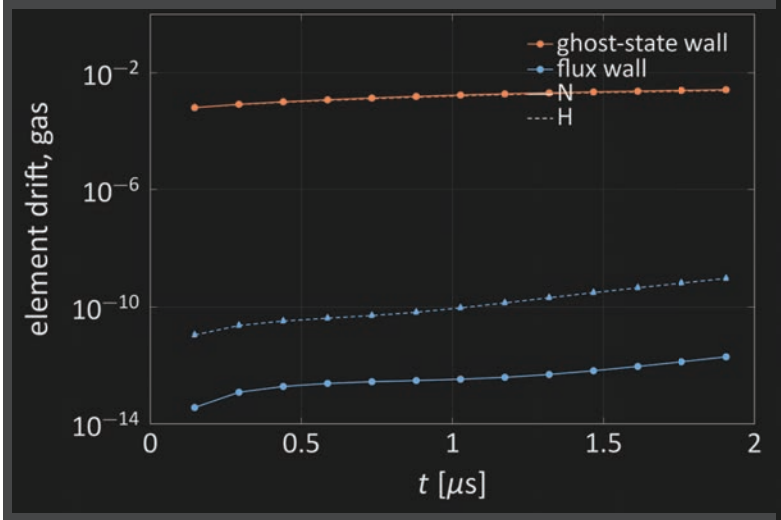

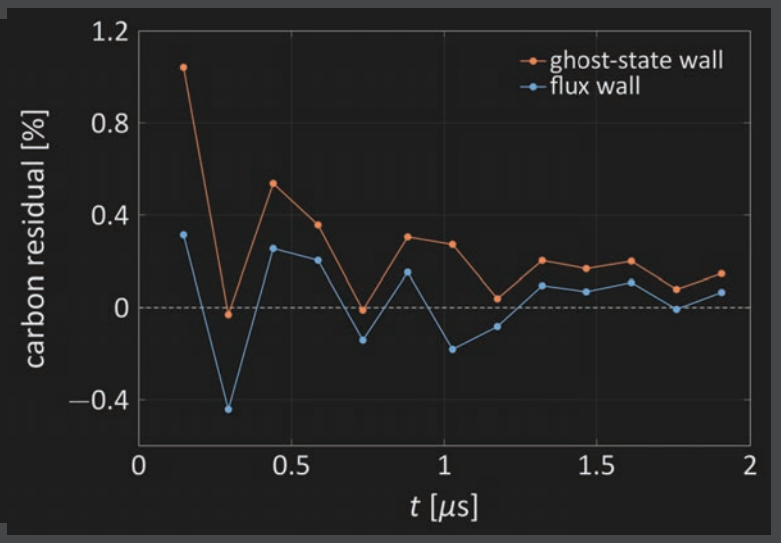

Flux wall: O, H, N to 10^{-12}-10^{-9}; carbon mean 1.6 x 10^{-4} | Ghost-state wall: loses 0.2-0.3%

Burning carbon sphere with a CO flame

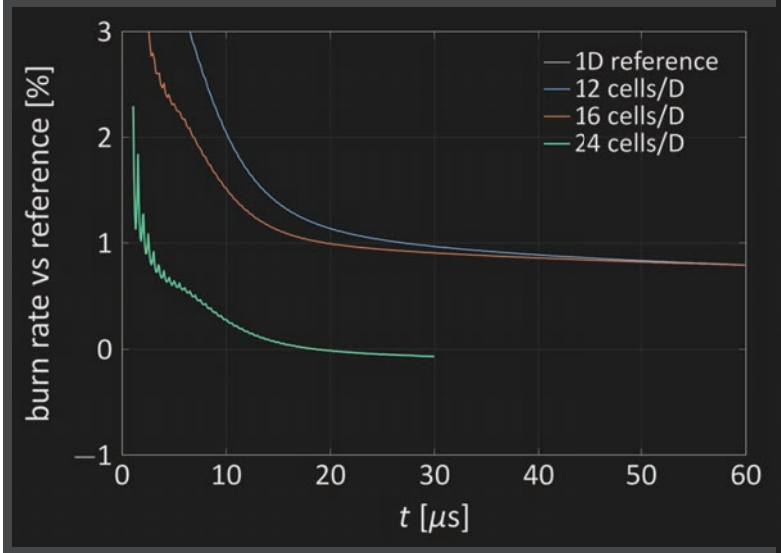

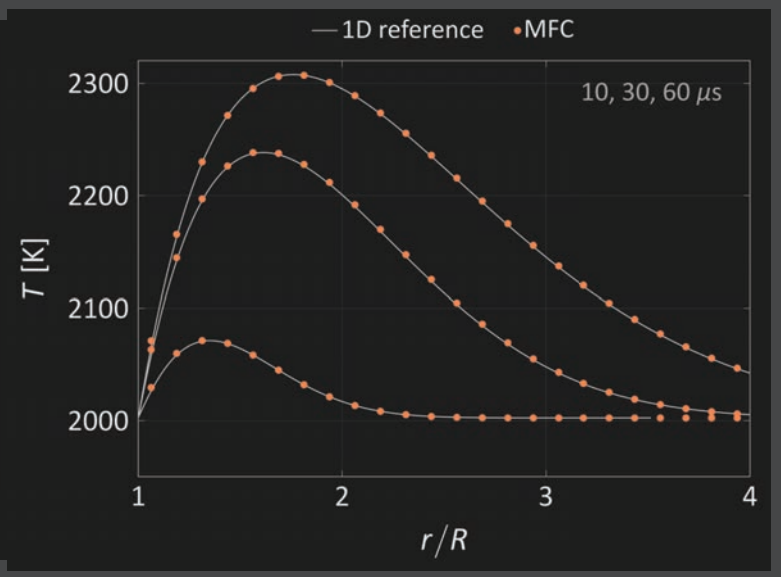

Burn rate within 0.3% of 1D spherical reference at 24 cells/D; flame T within 1.5 K

Detonations

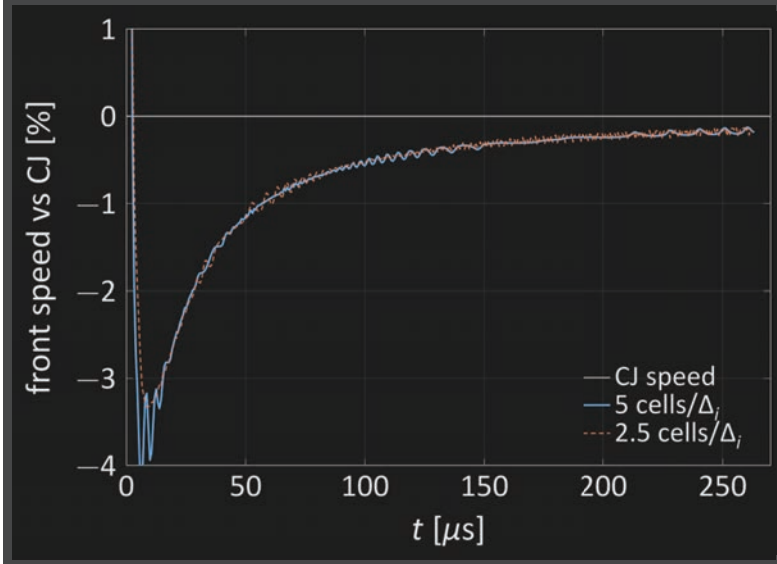

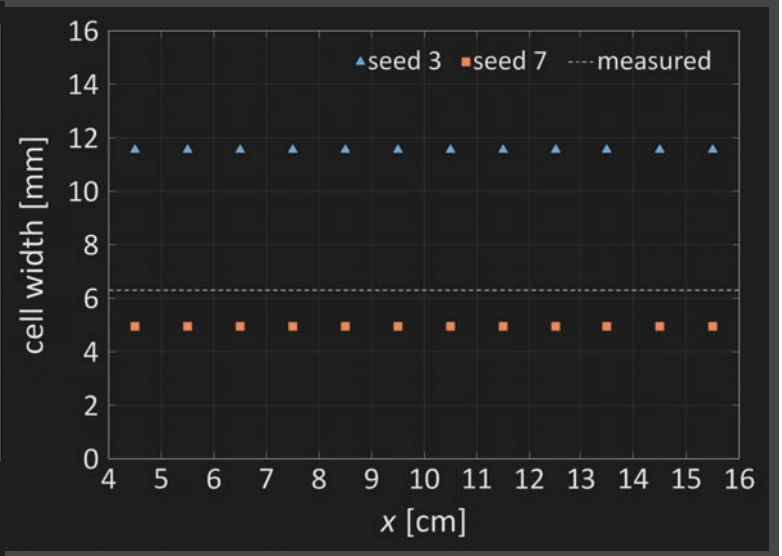

1D CJ speed: -0.2% | 2D cell width: 4.95-11.55 mm vs 6.3 mm measured

Detonation and flame structure

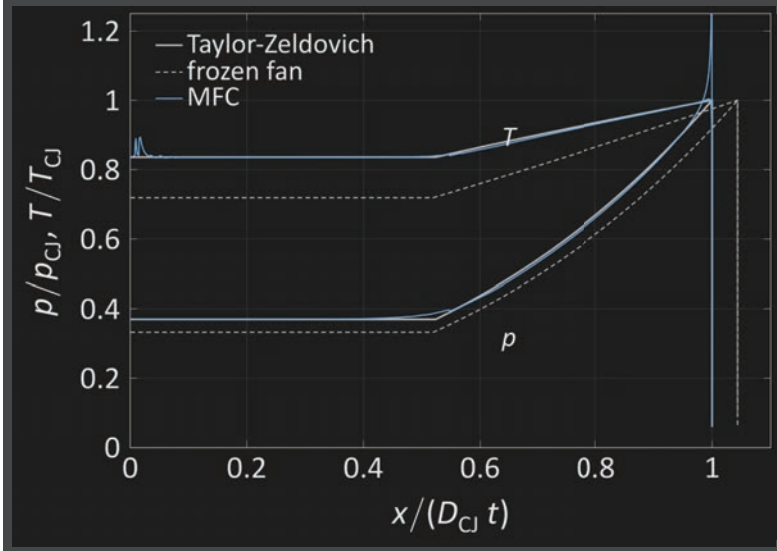

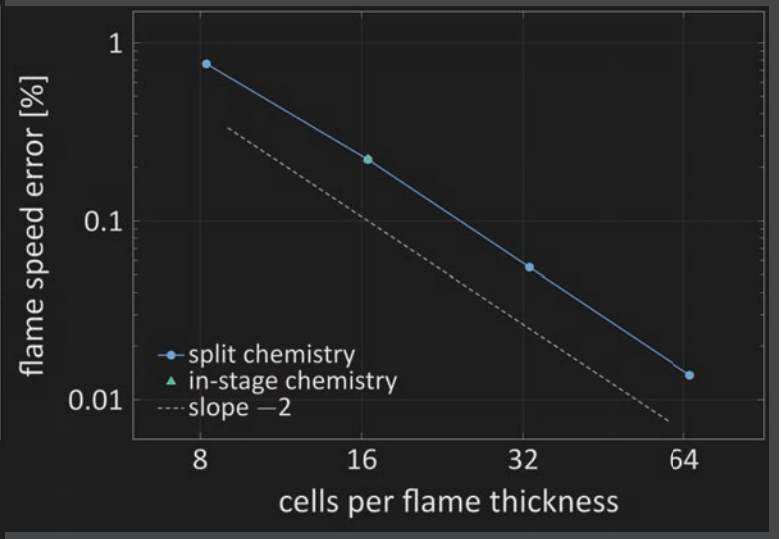

Burning carbon rods vs Makino experiments

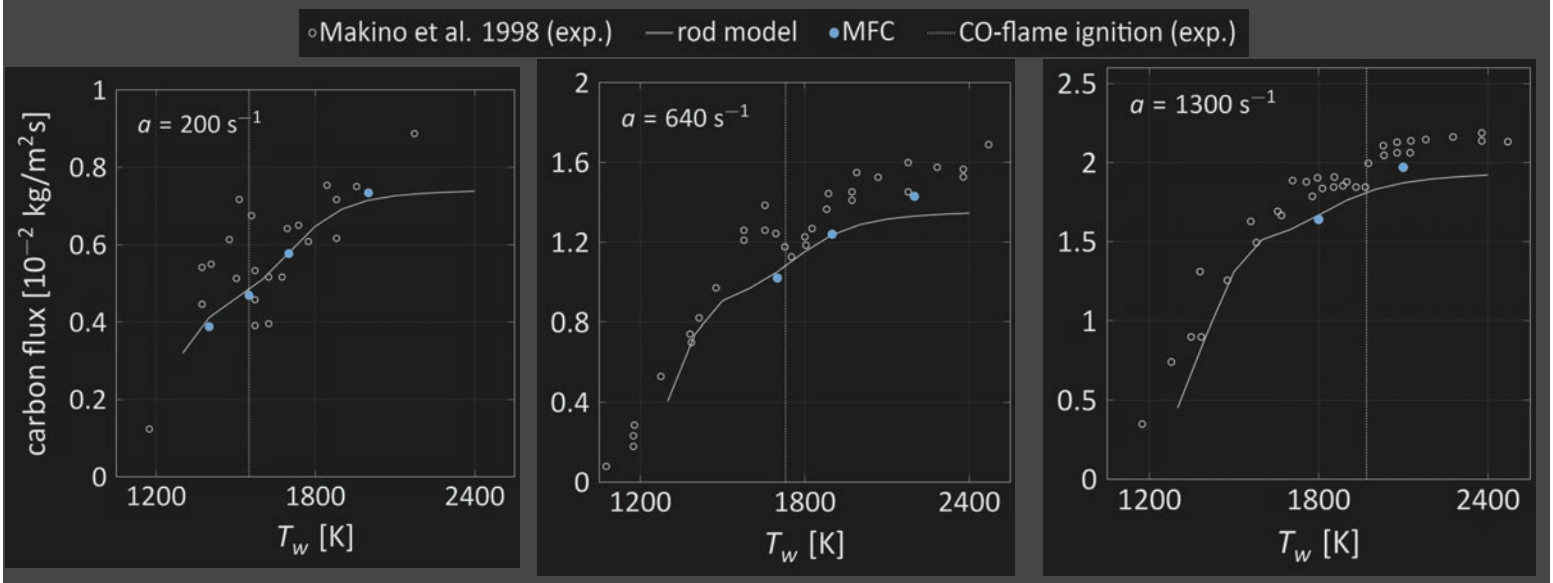

Three strain rates; MFC within 0.95-0.97 of 1D flame-surface model

64 burning particles

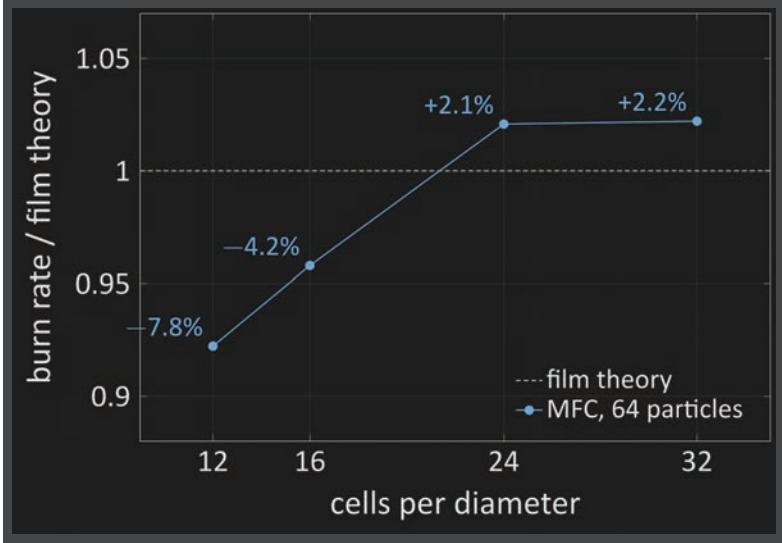

Periodic array, 2000 K air, 10 atm

Particles at 1800 K, Re 100

Burn rate vs film theory

+2% from 24 cells/D

24 to 32 cells/D: 0.06% change

Moving vs fixed bodies

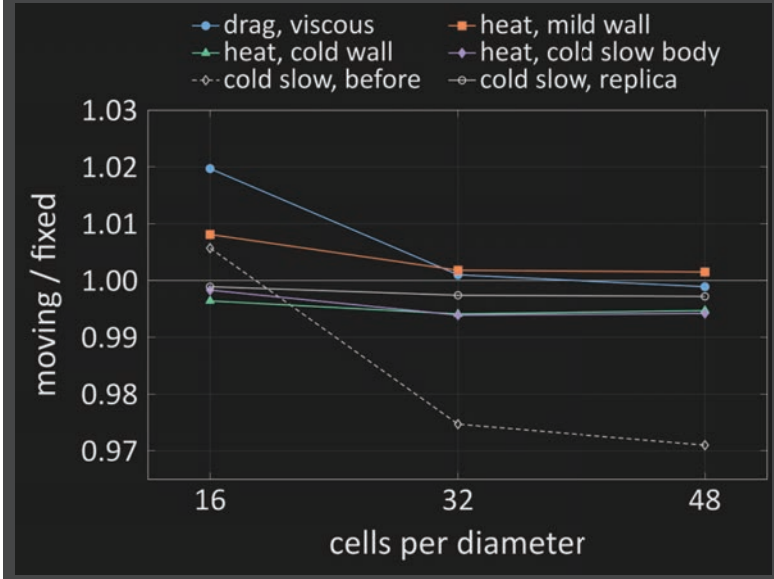

Cut-volume + GCL + moving flux wall

- Same heat and drag, moving or fixed
- Within 0.6% at 32-48 cells/D
- Cold slow body: 0.994 (was 0.971)
- Free stream exact to round-off

Cylinder in 1500 K air, Re 40; moving body over offset-averaged fixed twins. MFC campaign-v6 302fe209d

Euler-Lagrange vs Euler-Euler: cost

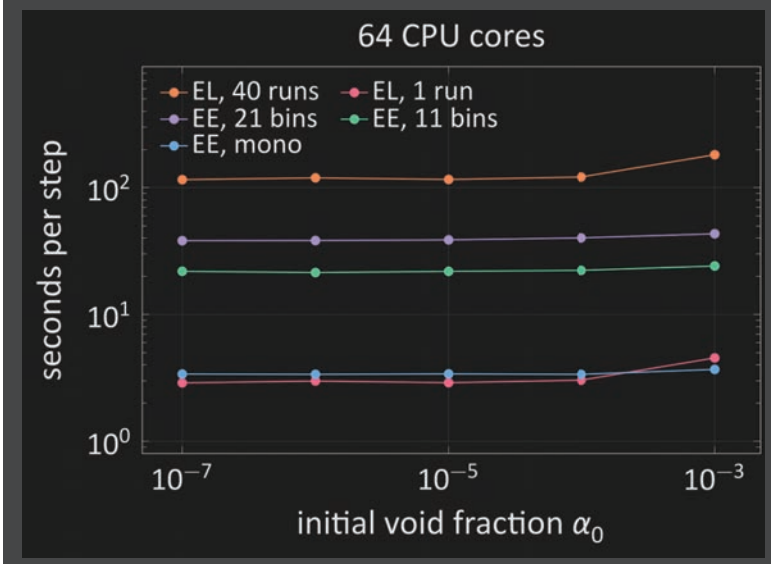

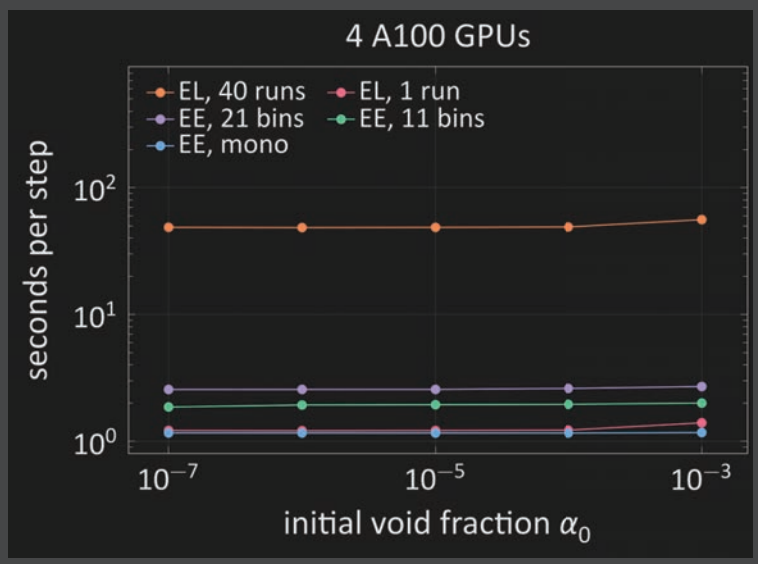

Vaca-Revelo et al., IJMF (2026)

Resolved particle collisions



Resolved bubble collapse in a cloud

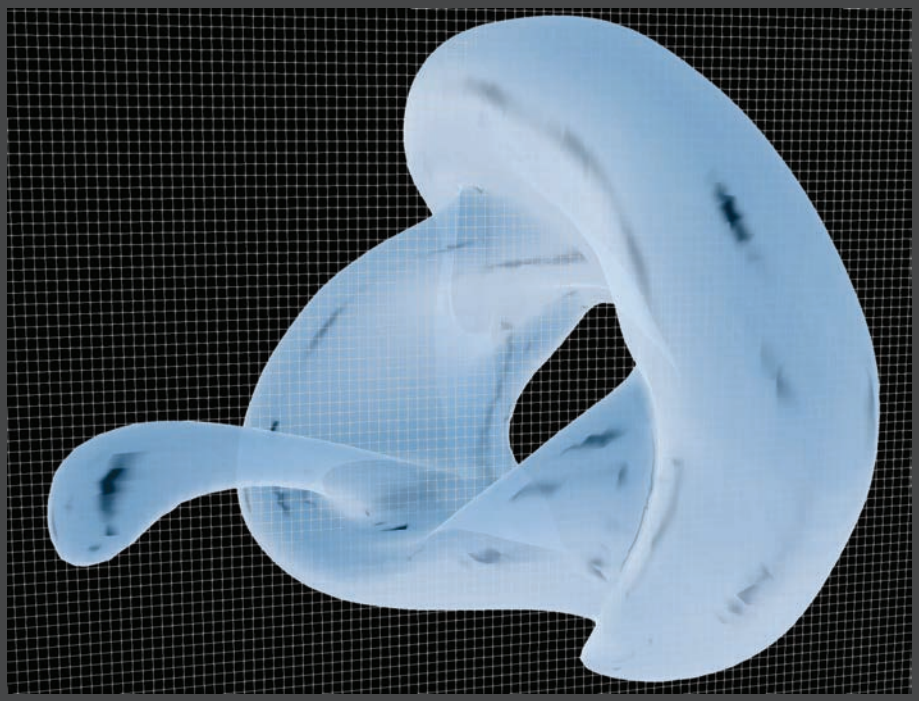

Close-up, one of 75 bubbles in a shock-driven cloud

MFC strong scaling: Frontier

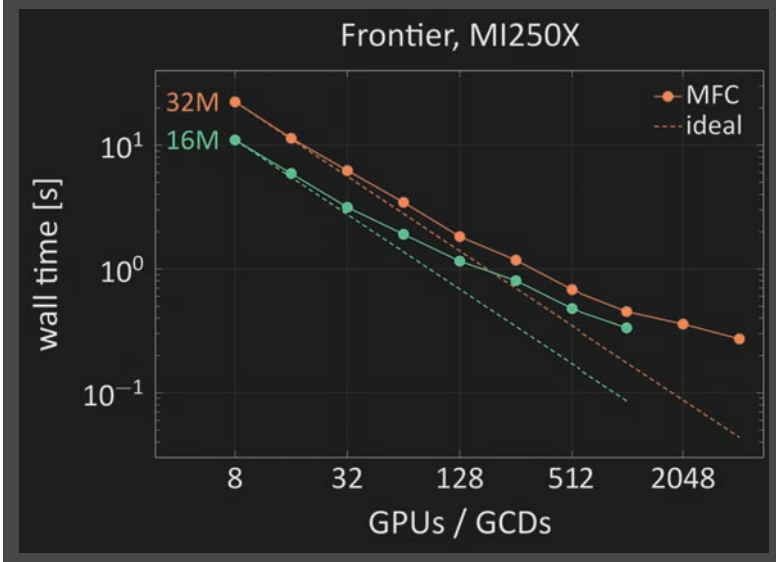

mflowcode.github.io | Wilfong et al., CPC (2026)

Second order at curved walls

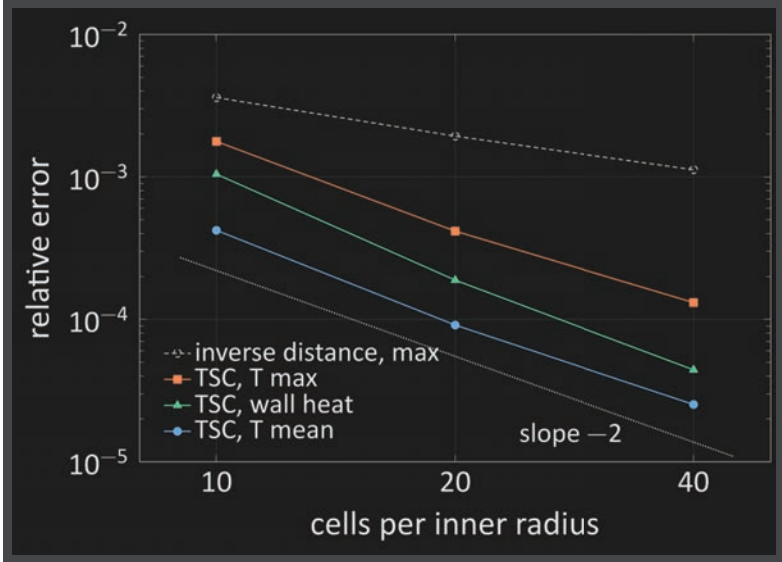

Annulus vs exact solution

- TSC image points: slope 2
- Inverse distance: slope 1

Wall heat and T field

- Second order, max norm

Euler-Lagrange particles: blast through a particle ring

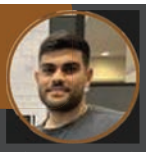
Thierry Daoud
U. Florida

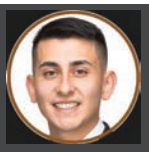
Jose Aguilar
U. Florida

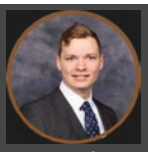
Ben Wilfong
Georgia Tech

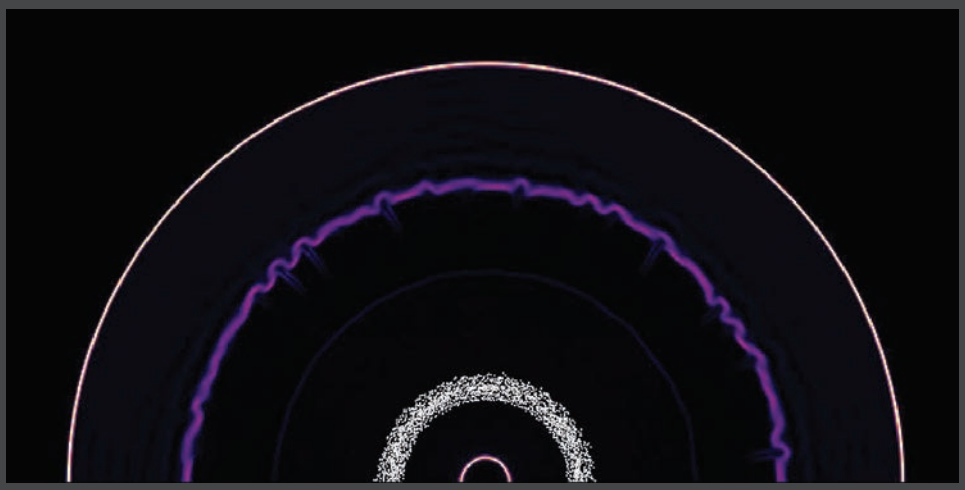

Mach 10 blast, 1,612 particles; color: |grad rho|/rho; t = 50 us

In MFC

Particles across MPI ranks

GPU = CPU, bitwise

Phase 1 (#1926)

Gidaspow, Parmar, Osnes drag

Pressure gradient, added mass

Two-way coupling

T. Daoud, J. Aguilar (UF); github.com/MFlowCode/MFC/pull/1926

AMR in MFC: verified

Conservation

Mass, momentum, energy to ~10^{-15}

Bitwise reproducible; rank-count independent to round-off

Accuracy

Second order in time, with regrid

Quartic fill: uniform-fine error

Minmod: 10^4-10^5x worse

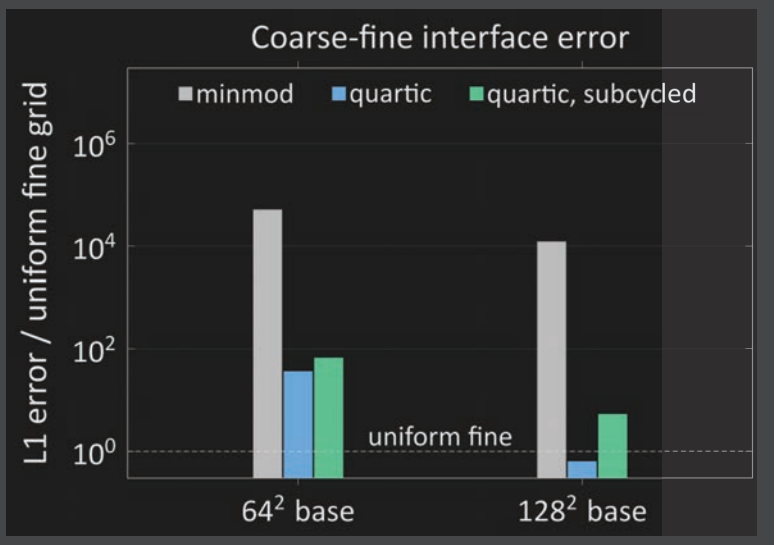

Euler-Lagrange: shock through a particle curtain

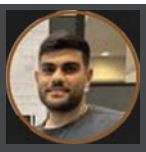
Thierry Daoud
U. Florida

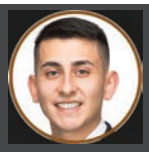
Jose Aguilar
U. Florida

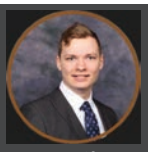
Ben Wilfong
Georgia Tech

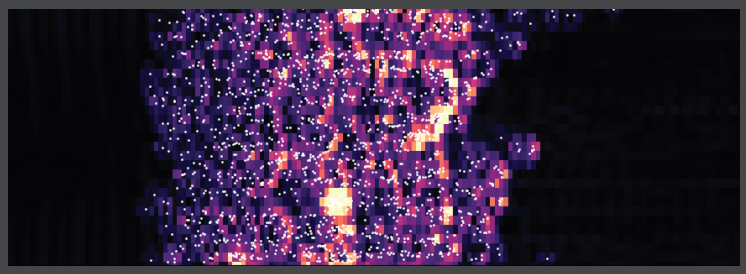

Planar shock, 1,978 particles, t = 500 us

Color: |grad rho|/rho; dots: particles

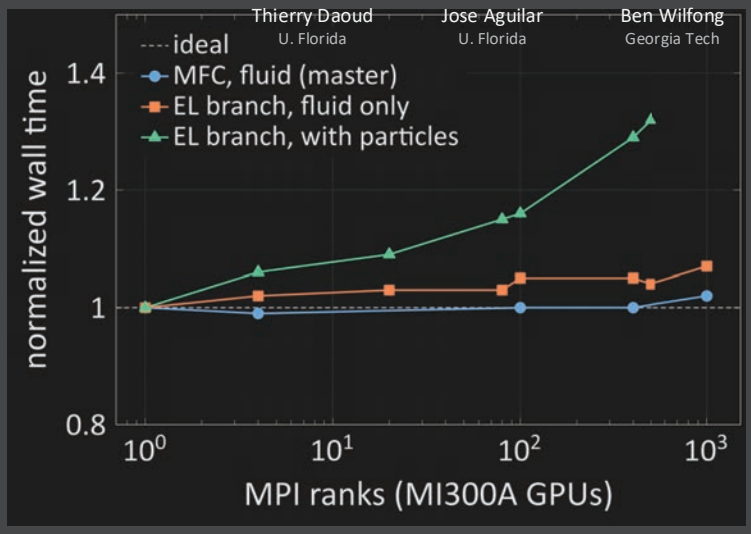

Weak scaling to 500 MI300A GPUs (Tuolumne)

Scaling digitized from PR #1926

Polydisperse closures

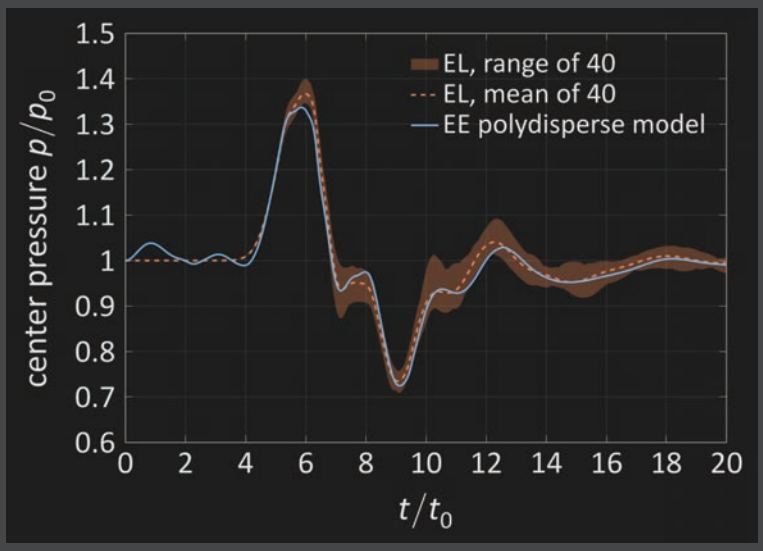

Polydisperse SGS model

Heat and mass transfer, many sizes

Radhakrishnan + Bryngelson, IJMF (2026)

Fourth-order HyQMOM

Multidimensional kinetic closures

Bryngelson, Fox, Laurent, JCP (2026)

Route to polydisperse EE in MFC

Reacting flow in MFC



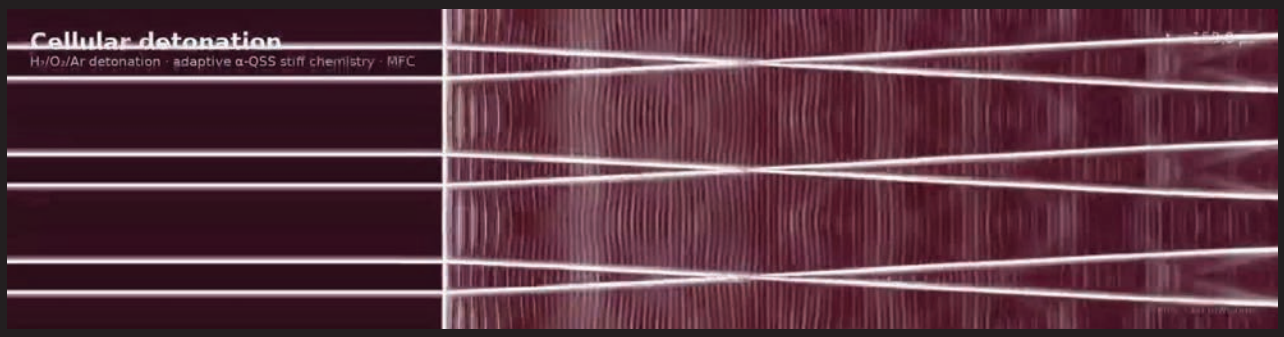

Reactive shock-bubble roll-up (H2-O2-Xe); cellular detonation (H2-O2-Ar); alpha-QSS chemistry

AMR load balancing: cost-aware regrid



Left: fixed spatial split. Right: blocks re-dealt by cost each regrid; 683 s vs 499 s wall clock

Euler-Lagrange: hemicylinder blast


Thierry Daoud
U. Florida


Jose Aguilar
U. Florida


Ben Wilfong
Georgia Tech

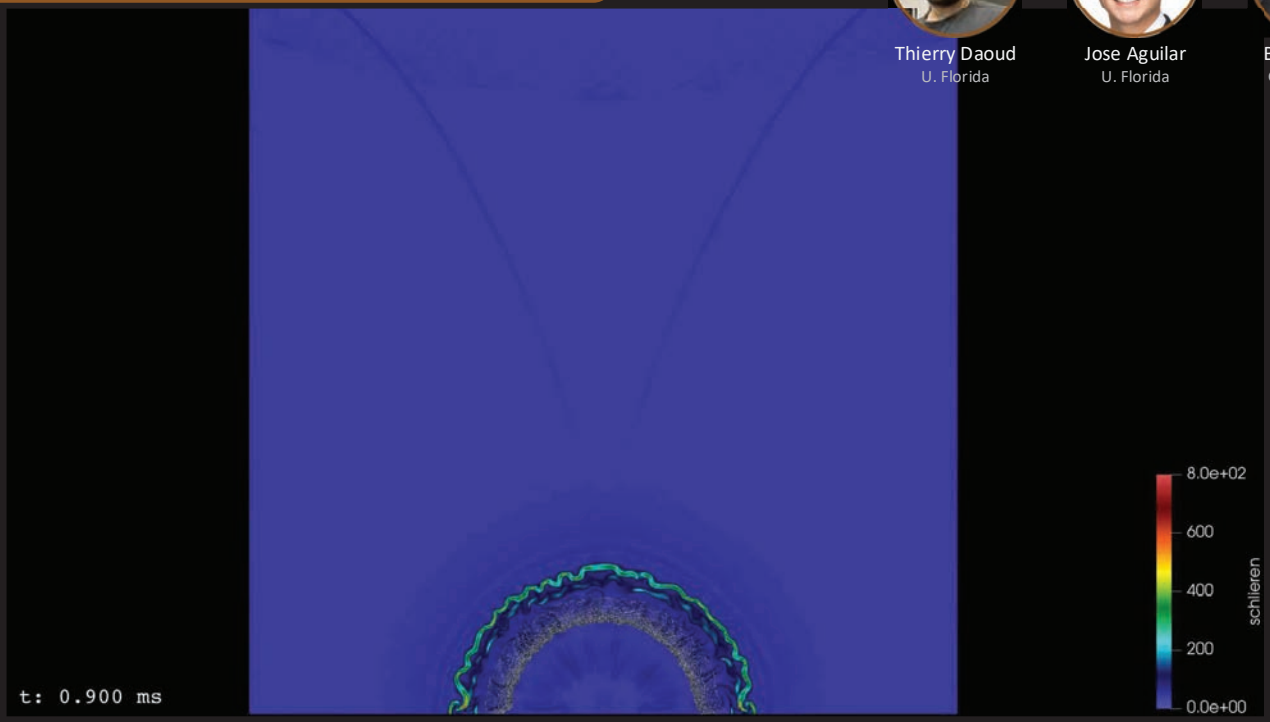

AMR: lock-step scaling and step time

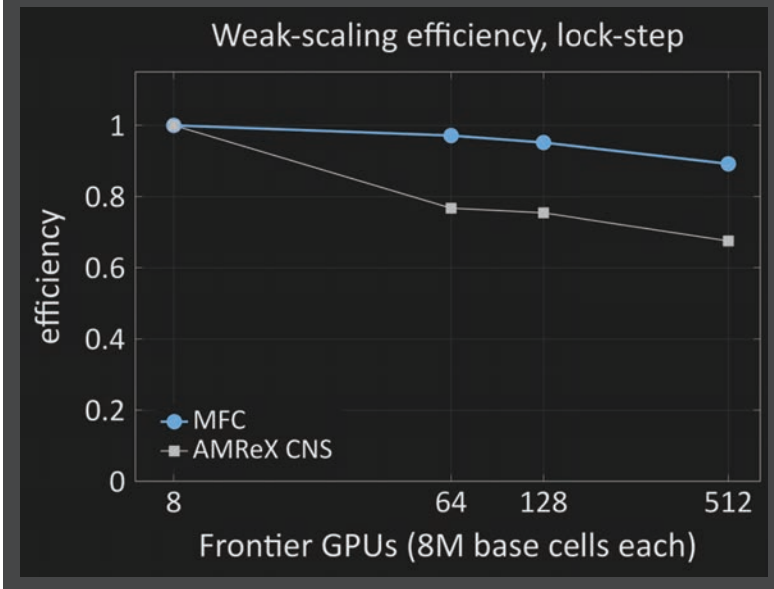

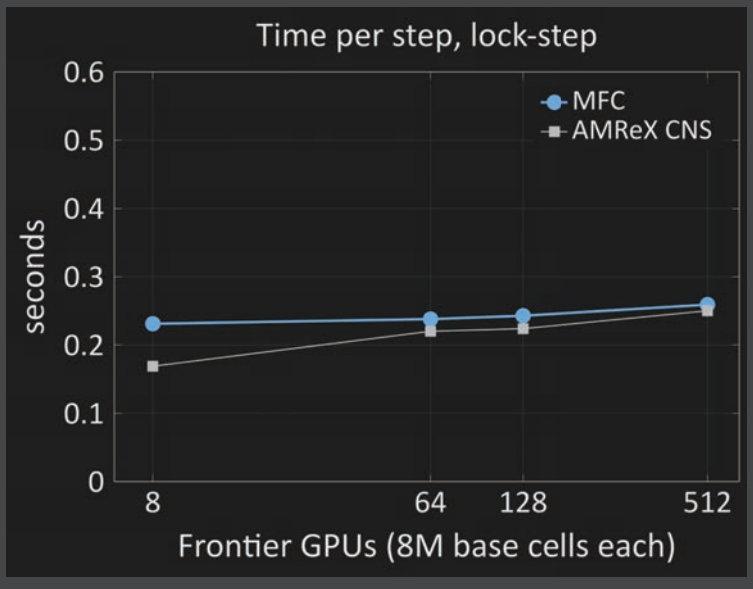

AMR subcycling: second order in time

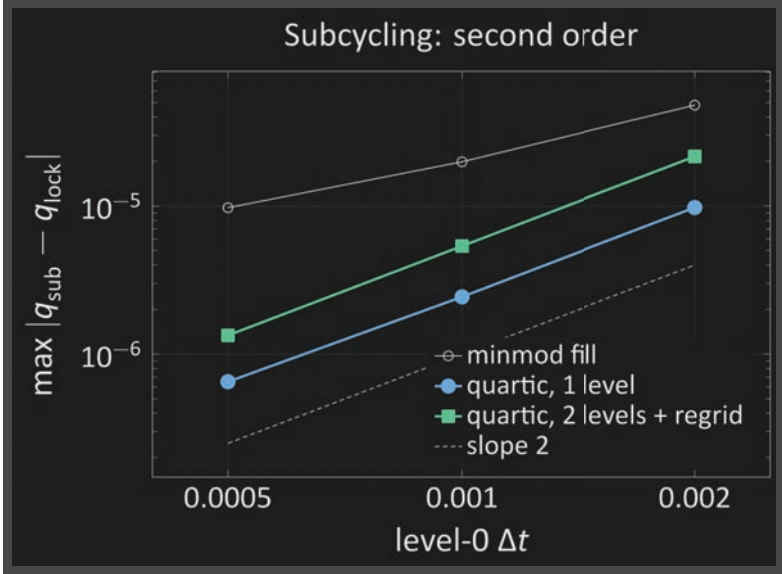

Quartic coarse-fine fill

- Second order, 1 and 2 levels
- Minmod fill: first order

AMR: as fast as PeleC or faster at scale

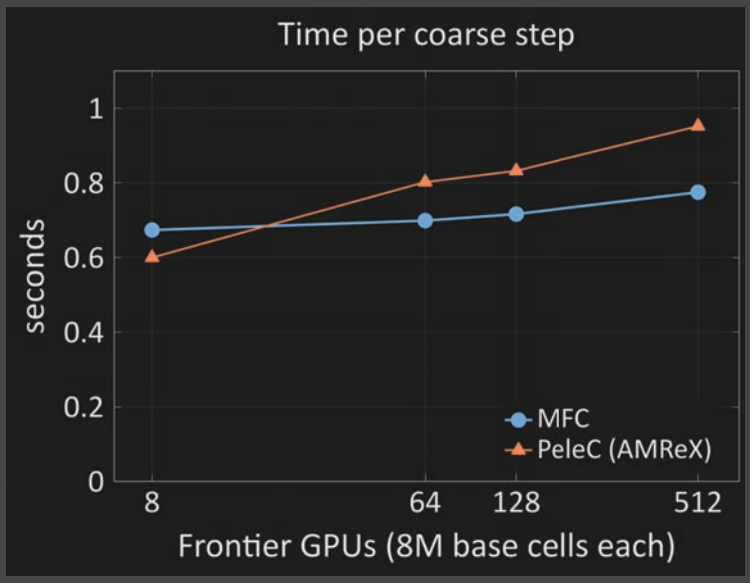

Note: MFC higher-order accurate here

Frontier, blast lattice, 8M base cells/GPU

AMR weak scaling incl. AMReX CNS

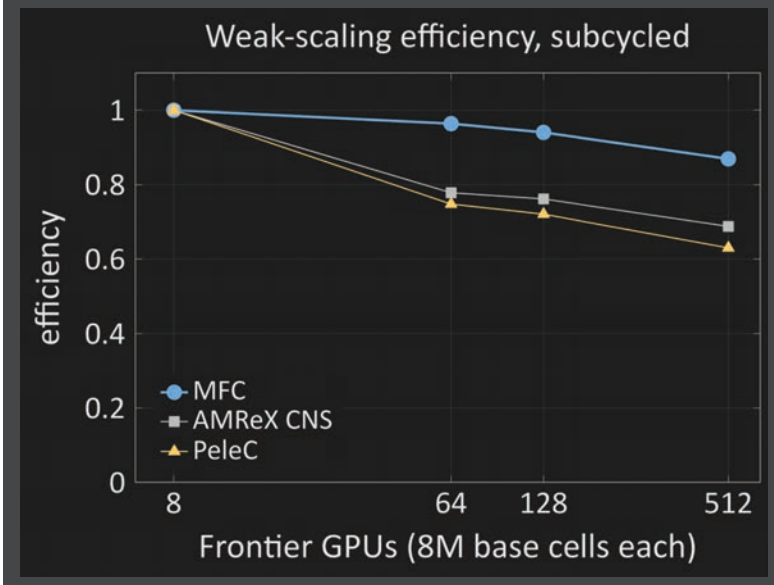

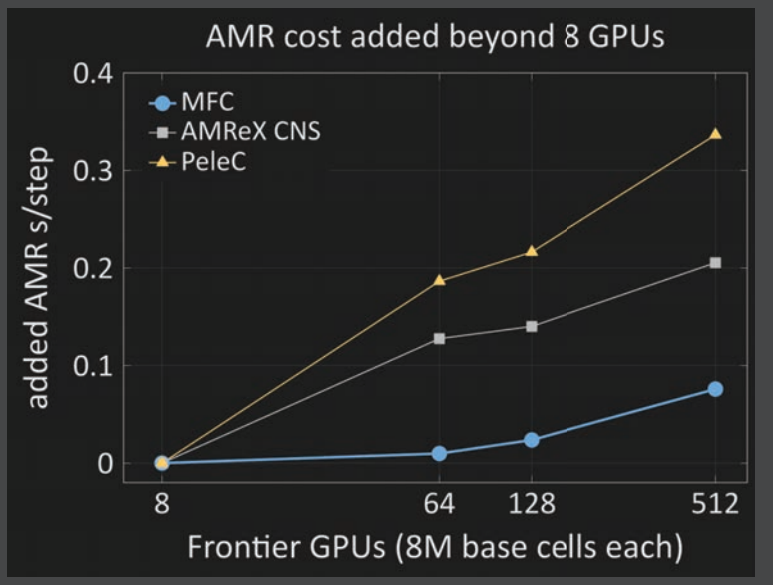

Blast lattice (AMReX/PeleC benchmark), subcycled, 512 GPUs: MFC 87%, AMReX 69%, PeleC 63%

Frontier, 8M base cells/GPU

AMR: time to solution

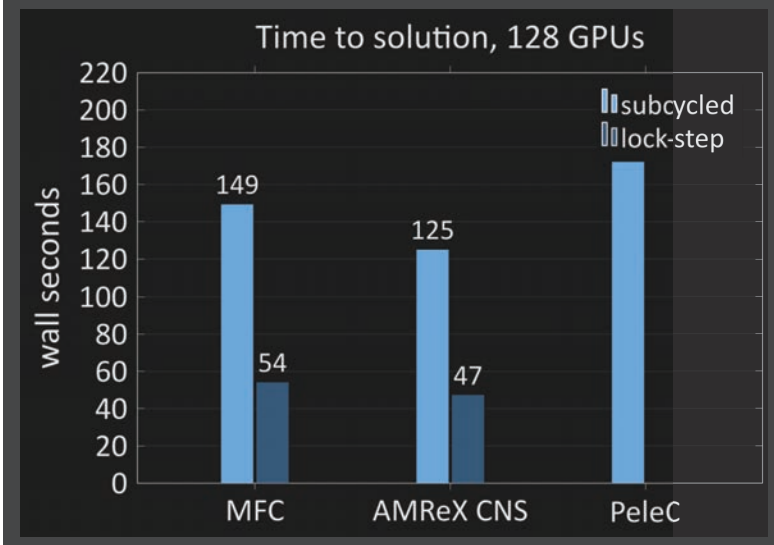

128 Frontier GPUs

Same physical time, blast lattice

MFC subcycled: 149 s

PeleC: 172 s

AMReX CNS: 125 s

AMR Load Balancing: Shock through eight free spheres

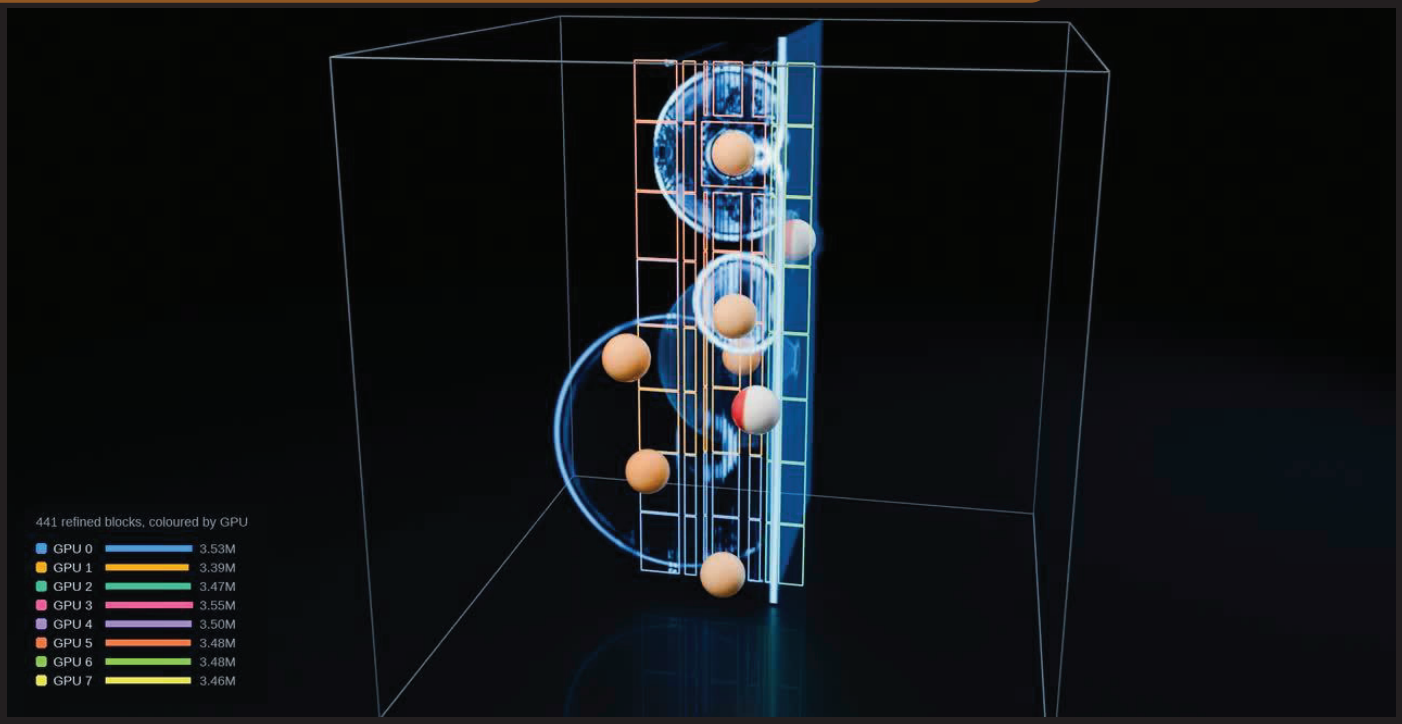

Mach 1.5, 8 force-driven spheres, 8 GPUs; refined blocks colored by GPU

AMR overhead and refined volume

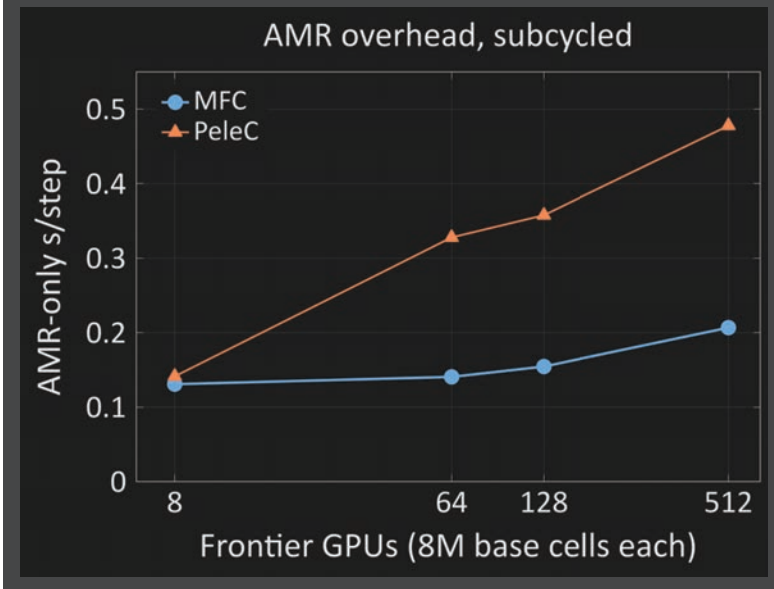

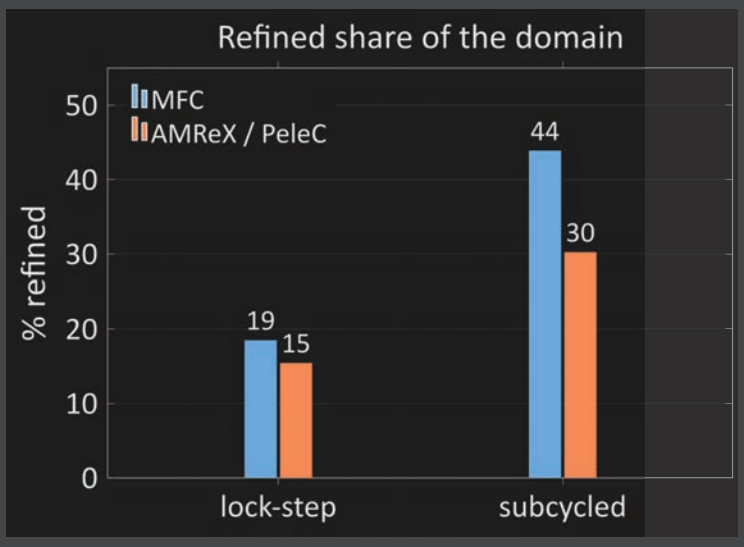

AMR-only = step time minus each code's own RHS; MFC refines ~1.3x more volume under the same tag rule

Low-Mach solver: formal convergence



Ben Wilfong
Georgia Tech

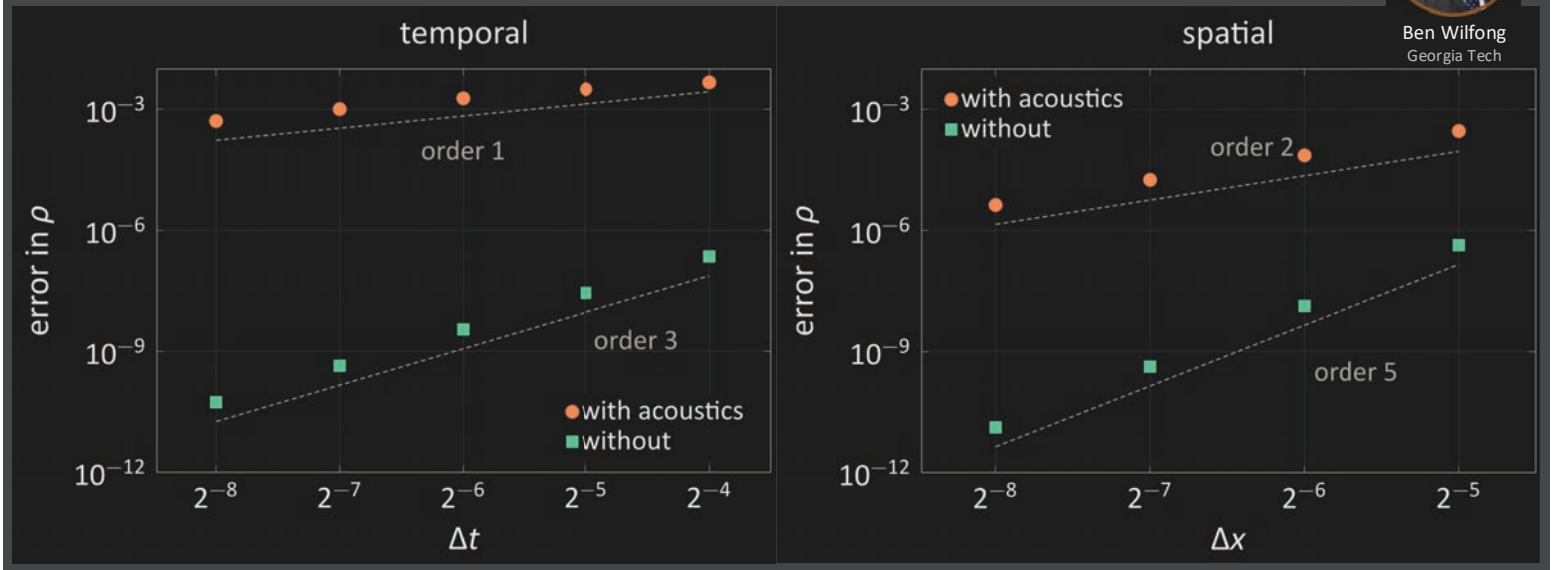

With acoustics: 1st order in time, 2nd in space; without: 3rd in time, 5th in space

Supersonic burning carbon sphere

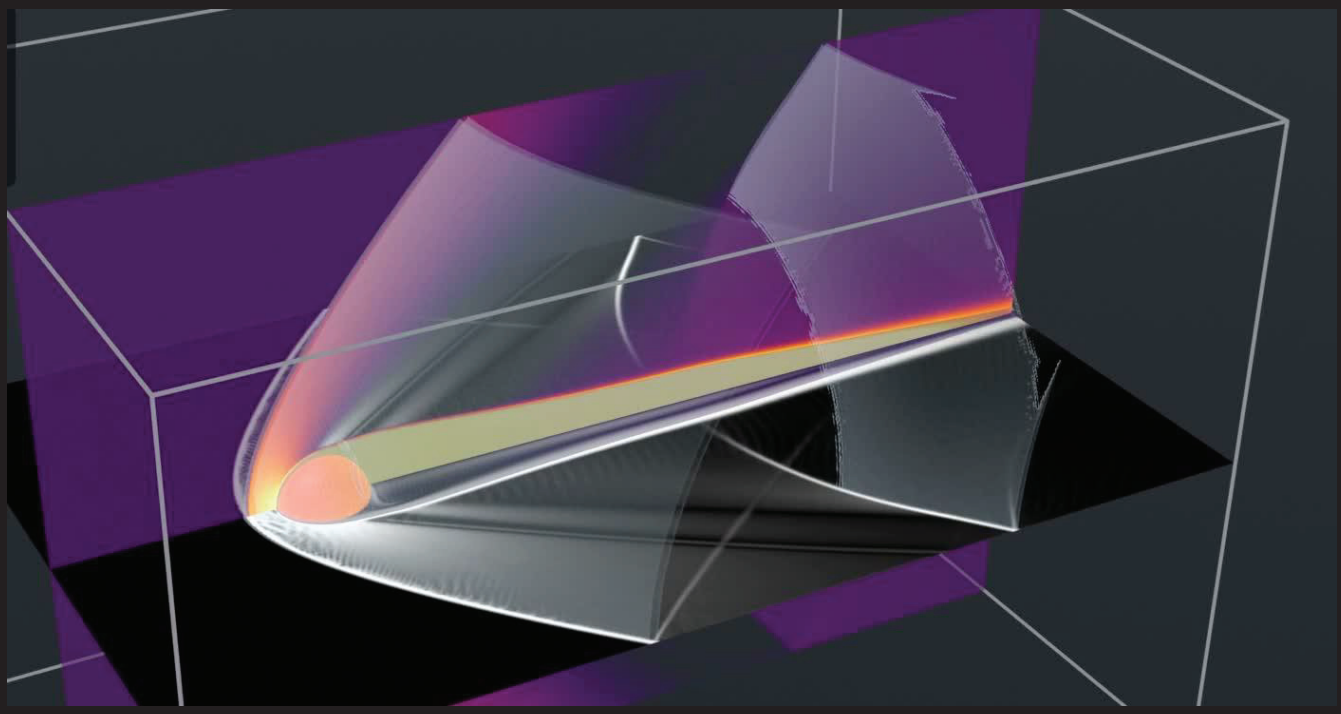

Mach 1.5, 49M cells: bow shock and CO plume

Polydisperse bubbles near a wall

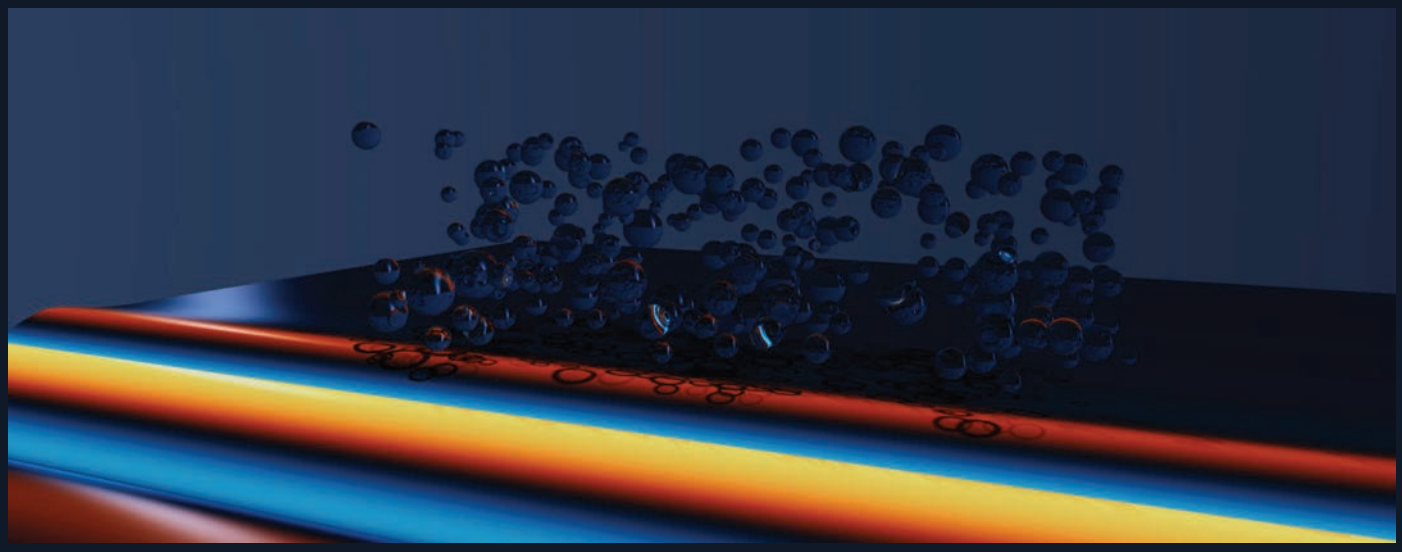

Polydisperse SGS model, heat + mass transfer: Radhakrishnan + Bryngelson, IJMF (2026)

4th-order HyQMOM kinetic closures: Bryngelson, Fox, Laurent, JCP (2026)

Euler-Euler: polydisperse bubble model

Mixture

$$\partial_t \mathbf{q} + \nabla \cdot \mathbf{F}(\mathbf{q}) = \mathbf{s}, \qquad p = (1-\alpha)\, p_l + \alpha \frac{\overline{R^3 p_{bw}} - \rho\, \overline{R^3 \dot{R}^2}}{\overline{R^3}}$$

$$\partial_t \alpha + \nabla \cdot (\alpha \mathbf{u}) = \alpha \nabla \cdot \mathbf{u} + 3\alpha\, \overline{R^2 \dot{R}} \,/\, \overline{R^3}$$

Bubble population: moments of $f(R, \dot{R}, R_0)$

$$\partial_t (n\, \mu_{lmn}) + \nabla \cdot (n\, \mu_{lmn} \mathbf{u}) = n\, g_{lmn}(R, \dot{R}, \ddot{R})$$

$$\mu_{lmn} \approx \sum_i w_i R_{0,i}^n \sum_{j,k} \hat{w}_{jk}\, \hat{R}_j^l\, \hat{\dot{R}}_k^m \qquad \text{(CHyQMOM nodes)}$$

Each node: Keller-Miksis + heat and mass transfer

$$R\ddot{R}\,(1 - \dot{R}/c) + \tfrac{3}{2}\dot{R}^2(1 - \dot{R}/3c) = \frac{p_{bw} - p_l}{\rho_l}(1 + \dot{R}/c) + \frac{R\,\dot{p}_{bw}}{\rho_l c}$$

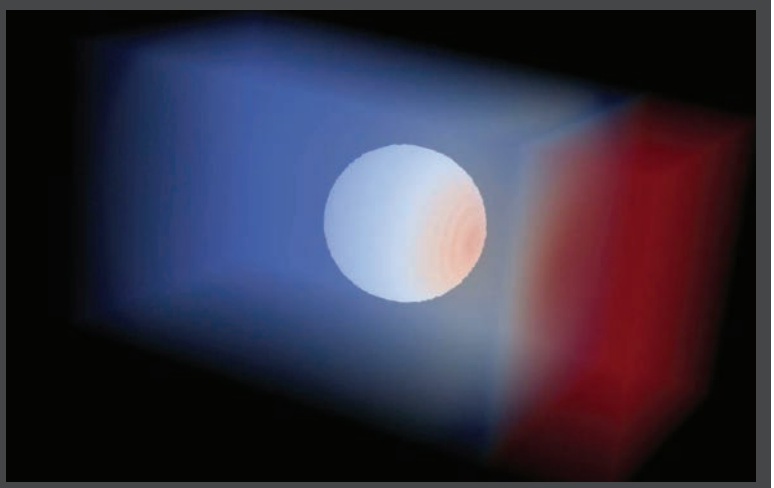

Shock through a bubble cloud, 3D, MFC on GPUs

Heat + mass transfer: Radhakrishnan + Bryngelson, IJMF (2026); 4th-order HyQMOM: Bryngelson, Fox, Laurent, JCP (2026)

Sixteen nodes, one Rabbit

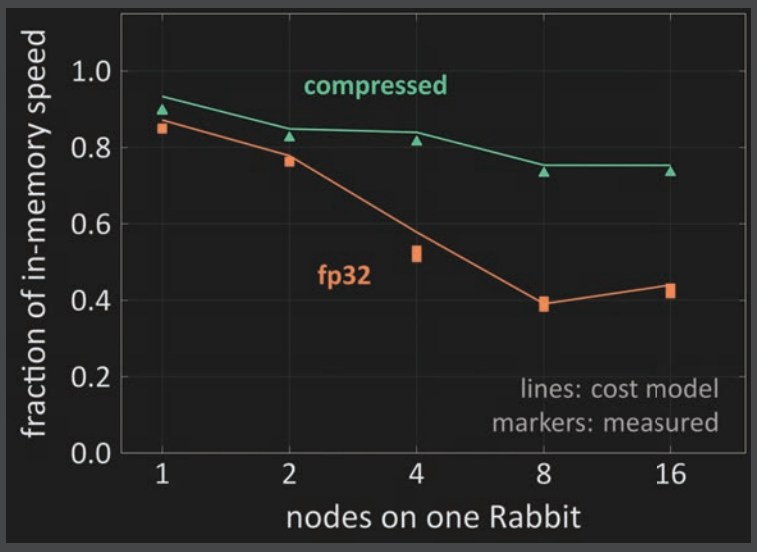

Shared flash

Compressed: 89% (1 node), 73% (16 nodes)

Accuracy

Bitwise equal to in-HBM run (fp32)

Rounding once per pass: spectrum 163x closer to fp64

Reacting-flow verification: all checks

Problem	Dim.	MFC result	Compared with
Flame hit by a vortex ring	3D	**2nd order**	H2O mass, flame area, 64^3 to 256^3
Flat H2-air flame	3D	**0.05%**	flame speed, Cantera
Burning carbon sphere, CO flame	3D	**0.3%**	burn rate, 1D spherical reference
64 burning particles	3D	**+2%**	film theory, grid converged
Element conservation, burning array	3D	**10^{-12}-10^{-9}**	O, H, N in a closed box
Burning carbon cylinder, CO flame	2D	**0.15%**	burn rate, 1D cylindrical reference
Carbon cylinder, kinetic to diffusion	2D	**0.98-1.00**	kinetic-diffusion theory
Burning carbon rods	2D	**0.95-0.97**	1D flame and surface model
Cellular detonation	2D	**4.95-11.55 mm**	6.3 mm cell width, measured
Detonation, H2-O2-Ar	1D	**-0.2%**	Chapman-Jouguet speed
Laminar H2-air flame	1D	**0.03%**	flame speed, Cantera
Reacting shock tube	1D	**-0.09 mm**	shock position, Martinez Ferrer et al.
Ignition	0D	**0.002%**	ignition delay, Cantera

Runs: Tuolumne, Tioga, Frontier. Chemistry: Cisneros-Garibay et al., Comput. Phys. Commun. 320, 109987 (2026)

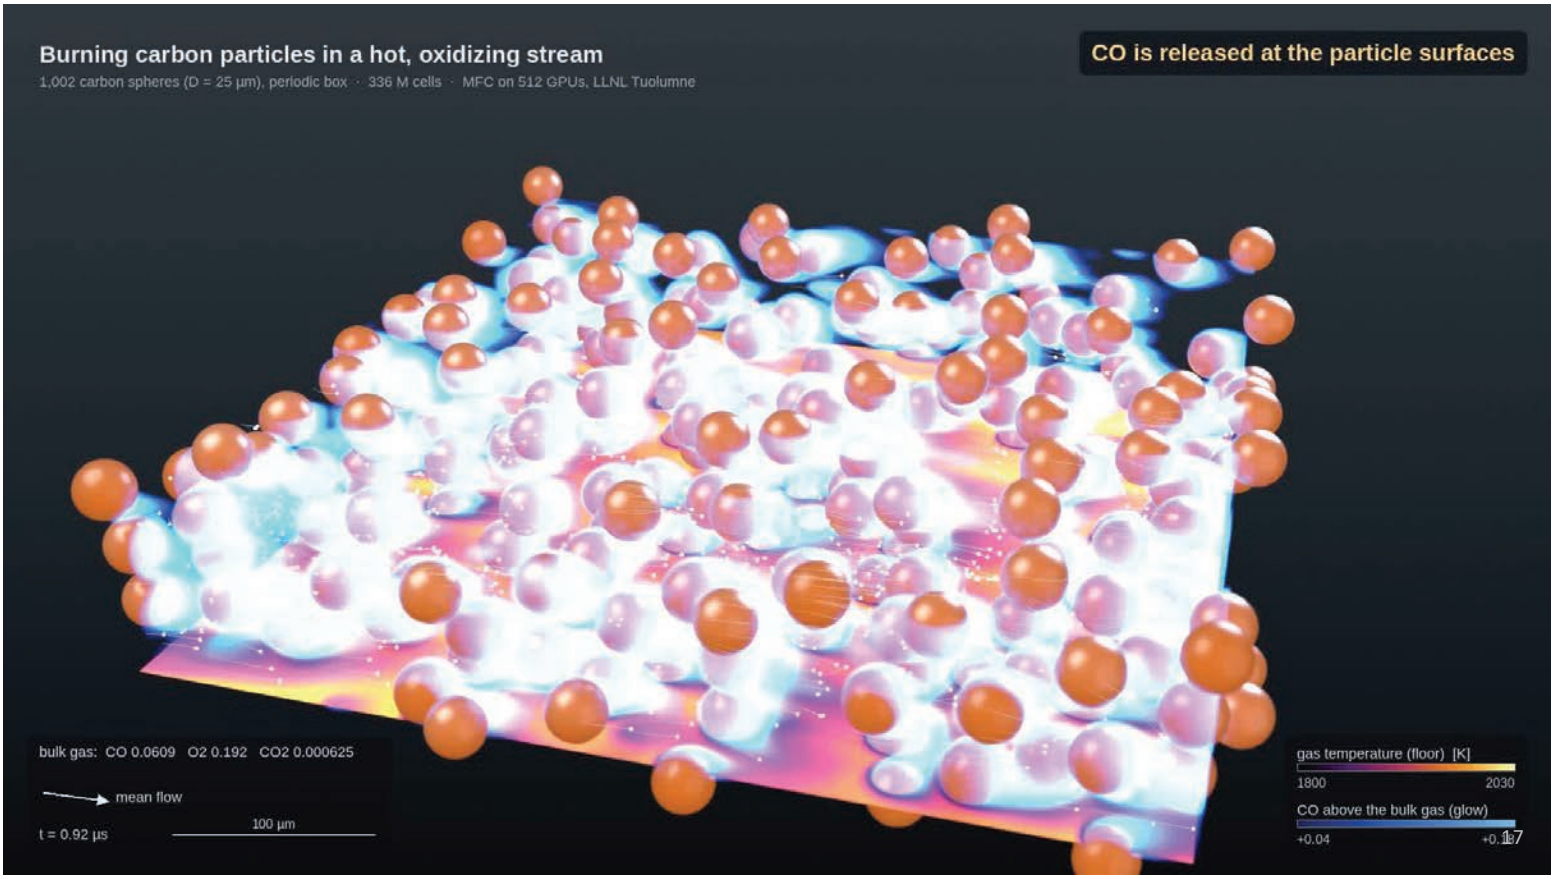

Reacting flow through a carbon particle bed

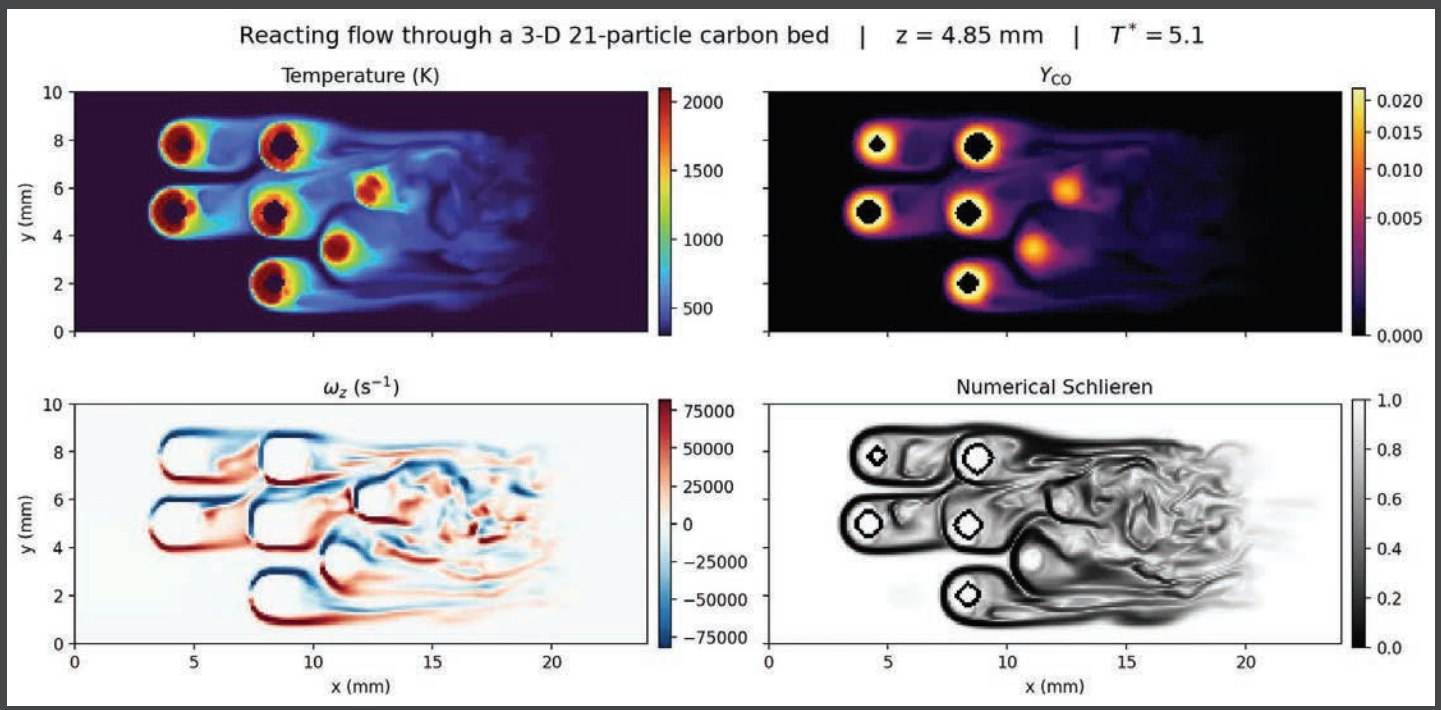

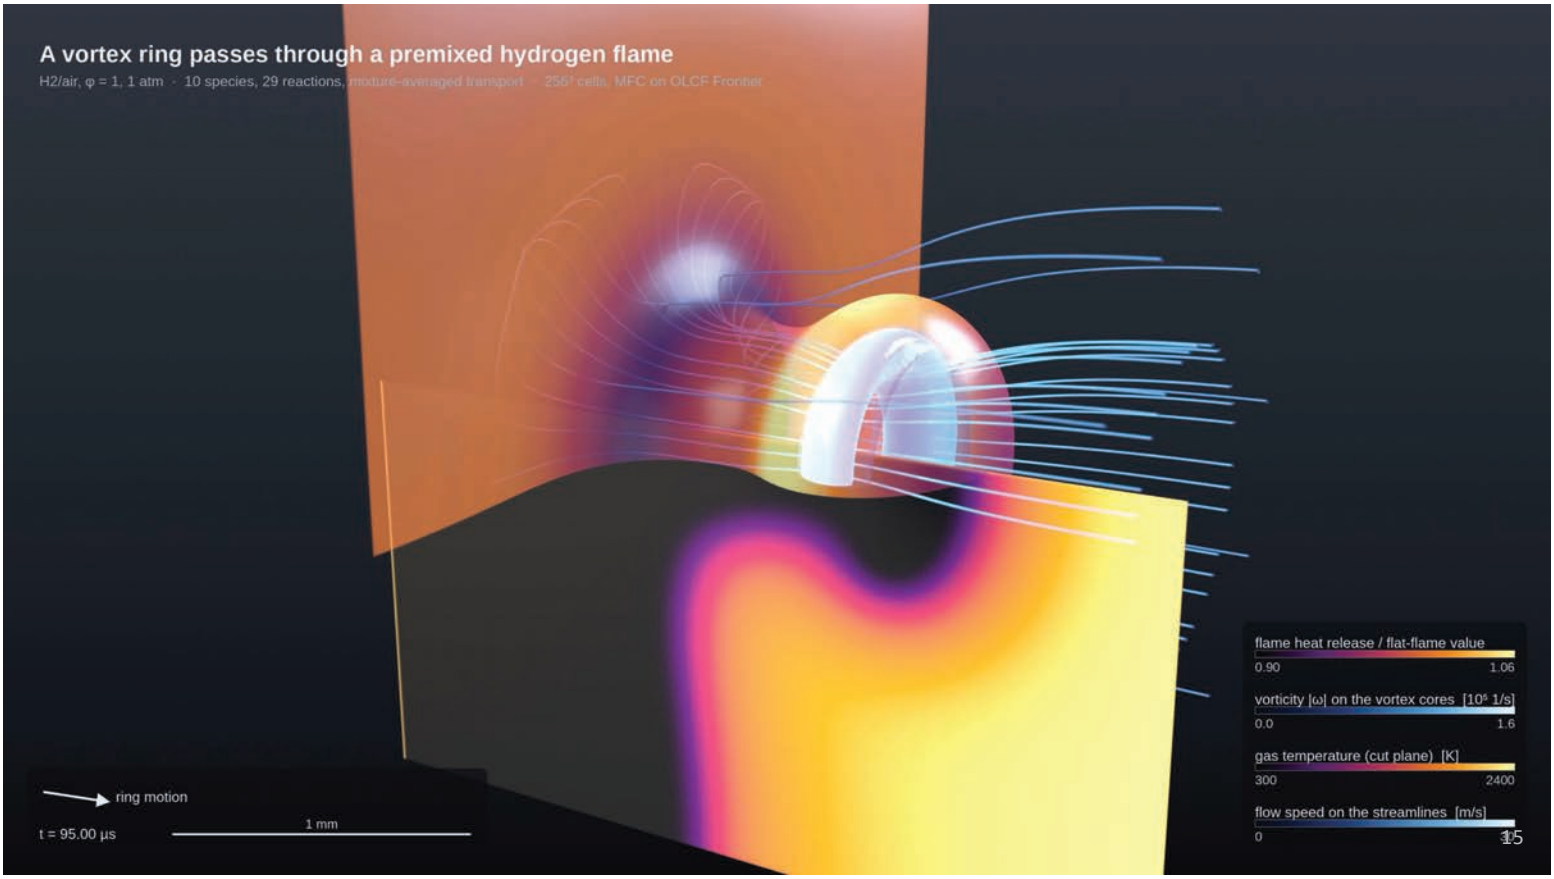

3D flame and vortex ring: 2nd-order

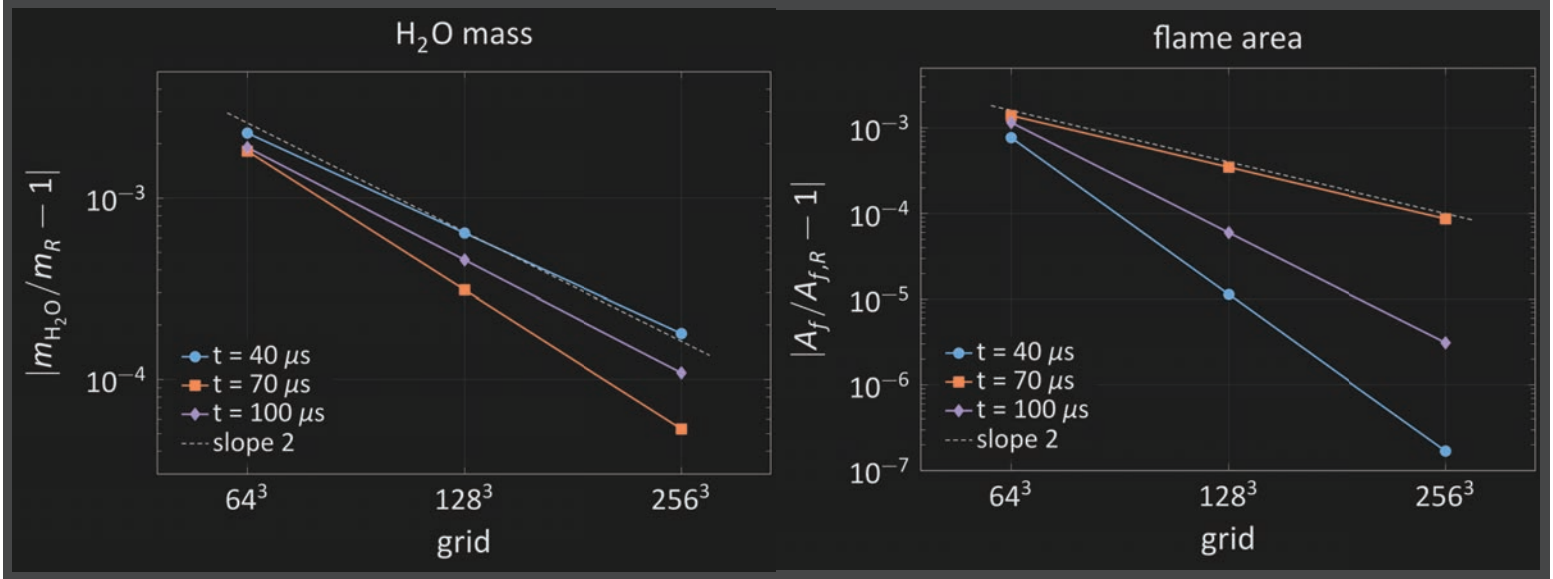

Error vs grid at three instants (t = 40, 70, 100 us)

Flat-flame run: flame speed within 0.05% of Cantera

CM³C: Center for Multiscale Modeling of Multiphase Combustion

T4. AI & ML Models

Joel B. Harley, Alina Zare, S. Balachandar

University of Florida

AI/ML Questions

Q1. Discuss how you have and will demonstrate the usefulness of this research in the context of your application?

Q2. How the results support and are integrated into the Center's simulation activities

Q3. How the results contribute to developing effective exascale computing in the broader community?

Q4. What specific AI/ML techniques are you using and developing?

Q5. How you have demonstrated the effectiveness of the research in the context of your application, including scaling, portability, and/or programmability?

Q6. What work is on the critical path and what challenges do you see in obtaining the required results?

Q7. How your Center plans to make (or has made) your software available to the Labs?

CM³C ML Team


S. Balachandar
(Florida)


Joel Harley
(Florida)


Alina Zare
(Florida)


James Fairbanks
(Florida)


Erica Lindbeck
(Florida)


BC Kim
(Florida)


Aayushma Aashu
(Florida)


Davin Anderson
(Florida)


Juan Diaz
(Florida)


Christian King
(Florida)


Cameron Lott
(Florida)


Isaly Tappan
(Florida)


Raul Valle
(Florida)

CM³C

Project Overview

This work

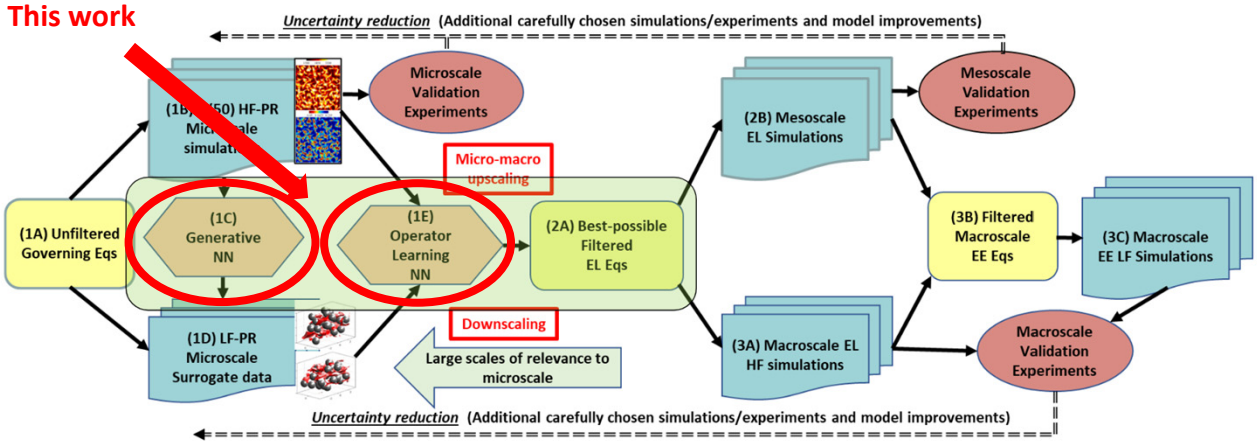

Microscale Mesoscale Macroscale

CM³C

Presentation Plan

- Motivation
- Background
- Overarching novel ML idea
- Data-Driven progress to date
- Physics-Informed modeling
- Next steps and other ideas

Motivation for ML Models

- Accurate simulations: Particle-resolved direct numerical simulations (PR-DNS)

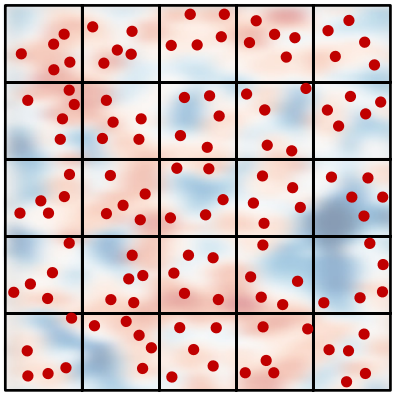

- PR-DNS across space are infeasibly expensive
- Requires computation of the full flow field and particle behavior

Motivation for ML Models

- A more computationally efficient option would be Euler-Lagrange Simulations

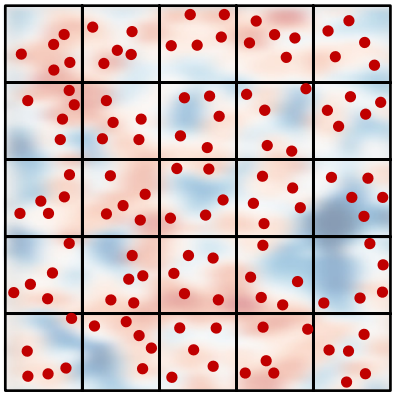

- Eulerian cell info: $\underline{\boldsymbol{u}}_f, \underline{\boldsymbol{T}}_f, \underline{\boldsymbol{a}}_f, \ldots, \nabla\underline{\boldsymbol{u}}_f, \nabla\underline{\boldsymbol{T}}_f, \ldots$
- **Lagrangian particles** much smaller than the grid
  - Particle info: $\underline{\boldsymbol{x}}_i, \underline{\boldsymbol{v}}_i, \boldsymbol{d}_i, \boldsymbol{m}_i, \ldots$
- Much faster to run than PR-DNS

Motivation for ML Models

- **Common Solution:** Filtering the governing equations or ensemble averaging

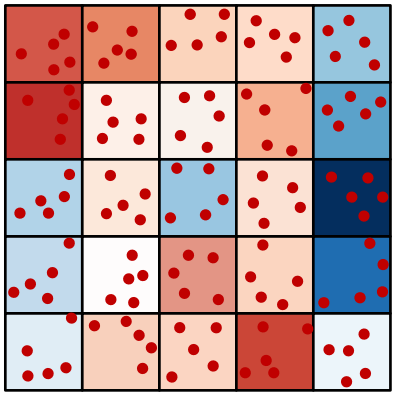

- Unrealistic approximation
- Based on homogenizing statistics rather than a deterministic solutions

Mesoscale Simulations – Closure Terms

- Exact filtered (phase-averaged) equations:

$$\frac{\partial}{\partial t}\left(\phi_g \overline{\rho}_g\right) + \nabla \cdot \left(\phi_g \overline{\rho}_g \tilde{\mathbf{u}}_g\right) = 0$$

$$\frac{\partial}{\partial t}\left(\phi_g \overline{\rho}_g \tilde{\mathbf{u}}_g\right) + \nabla \cdot \left(\phi_g \overline{\rho}_g \tilde{\mathbf{u}}_g \otimes \tilde{\mathbf{u}}_g\right) = -\phi_g \nabla \overline{p}_g + \boxed{\phi_g \nabla \cdot \overline{\tau}_g} - \boxed{\nabla \cdot \left(\phi_g \mathbf{R}_{sg}\right)} - \boxed{\mathbf{F}'_g}$$

$$\frac{\partial}{\partial t}\left(\phi_g \overline{\rho}_g \tilde{E}_g\right) + \nabla \cdot \left(\phi_g \overline{\rho}_g \tilde{\mathbf{u}}_g \tilde{E}_g\right) = -\nabla \cdot \left(\phi_g \overline{p}_g \tilde{\mathbf{u}}_g\right) + \boxed{\nabla \cdot \left(\phi_g \overline{\tau}_g \cdot \tilde{\mathbf{u}}_g\right)}$$

$$- \overline{p}_g \nabla \cdot \left(\phi_p \mathbf{u}_p\right) + \boxed{\overline{\tau}_g : \nabla\left(\phi_p \mathbf{u}_p\right) - \phi_g \nabla \cdot \overline{\mathbf{q}}_g} - \boxed{\left(W'_q + Q'_q\right)} - \boxed{\nabla \cdot \left(\phi_g \mathbf{Q}_{sg}\right)}$$

Viscous/Diffusive | Surface Integrals | Closure Terms

- Filtered quantities to be computed using particle-resolved simulations (PRS):

- $\mathbf{R}_{sg} = \overline{\rho \mathbf{u} \otimes \mathbf{u}} - \frac{1}{\overline{\rho}} \overline{\rho \mathbf{u}} \otimes \overline{\rho \mathbf{u}}$
- $\mathbf{Q}_T = \gamma C_v \left(\overline{\rho \mathbf{u} T} - \tilde{T}\,\overline{\rho \mathbf{u}}\right)$,
- $\mathbf{E}_{ku} = \frac{1}{2}\left[\overline{\rho \mathbf{u}(\mathbf{u} \cdot \mathbf{u})} - \frac{\overline{\rho \mathbf{u}}}{\overline{\rho}} \overline{\rho(\mathbf{u} \cdot \mathbf{u})}\right]$
- $\mathbf{W}_{\tau u} = \frac{1}{\overline{\rho}}\left(\overline{\rho \boldsymbol{\tau} \cdot \mathbf{u}} - \overline{\boldsymbol{\tau}} \cdot \overline{\rho \mathbf{u}}\right)$
- $\mathbf{R}^g_{\mu,sg} = \overline{\tau} - \frac{2}{3}\mu_g (\nabla \cdot \overline{\mathbf{u}}_m)\mathbf{I} + \mu_g \left[(\nabla \overline{\mathbf{u}}_m) + (\nabla \overline{\mathbf{u}}_m)^{\mathsf{T}}\right]$
- $\mathbf{R}^g_{\lambda,sg} = \overline{\mathbf{q}} - \lambda_g \nabla(\tilde{T})$

- Individual components of the filtered sorted into scalar, vector, and second-rank tensor:

- **Scalar:** $\overline{\phi}_p$, $\overline{\rho}$, $\overline{\rho(\mathbf{u} \cdot \mathbf{u})}$
- **Vector:** $\overline{\mathbf{u}}_p$, $\overline{\rho \mathbf{u}}$, $\overline{\rho \mathbf{u} T}$, $\overline{\rho \mathbf{u}(\mathbf{u} \cdot \mathbf{u})}$, $\overline{\rho \boldsymbol{\tau} \cdot \mathbf{u}}$, $\overline{\mathbf{q}}$
- **Tensor:** $\overline{\rho \mathbf{u} \otimes \mathbf{u}}$, $\overline{\boldsymbol{\tau}}$

CM³C

Motivation for ML Models

- Need to **learn** closure terms for Euler-Lagrange Simulations

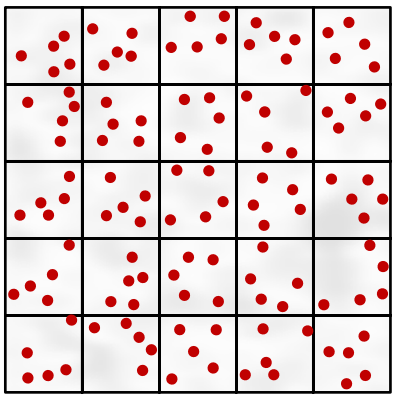

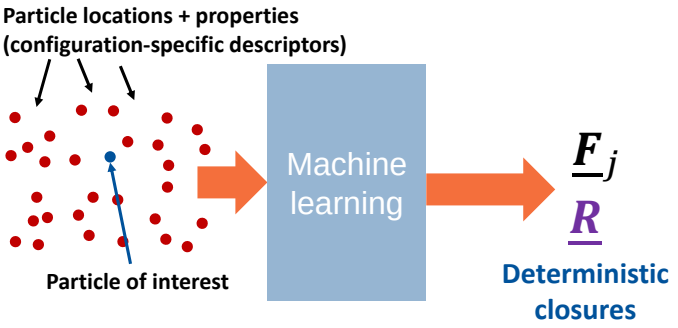

CM³C

Motivation for ML Models

- **Problem:** Particles change flow field. Flow field affects particles.

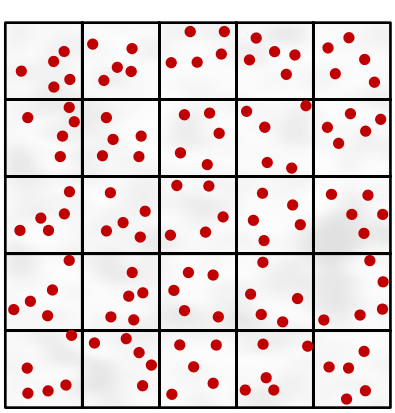

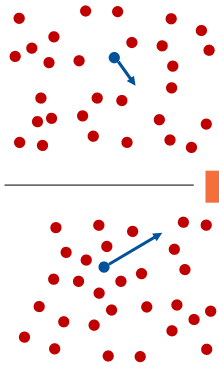

QoI: Weaker force
lower Reynolds stress

QoI: Larger force
higher Reynolds stress

Machine learning

$\underline{\boldsymbol{F}}_j$

$\underline{\boldsymbol{R}}$

Deterministic closures

Same mesoscale particle & flow descriptors
but different microscale neighborhoods
⇒ Different Quantity of Interest (QoI)

Motivation for ML Models

- **Problem:** Particles change flow field. Flow field affects particles.

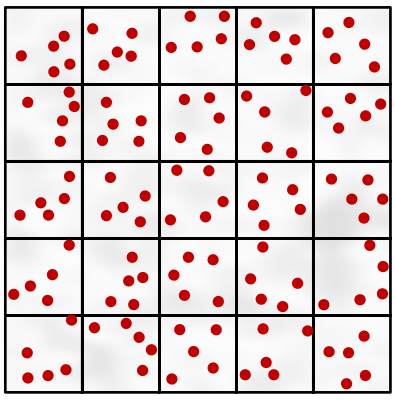

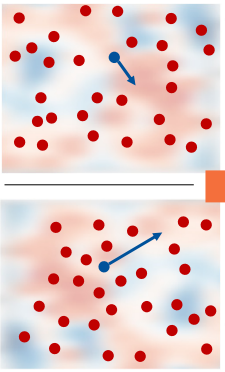

QoI: Weaker force
lower Reynolds stress

QoI: Larger force
higher Reynolds stress

Machine learning

$\underline{\boldsymbol{F}}_j$

$\underline{\boldsymbol{R}}$

Deterministic closures

Same mesoscale particle & flow descriptors
but different microscale neighborhoods
⇒ Different Quantity of Interest (QoI)

CM³C

Motivation for ML Models

- **Problem:** Particles change flow field. Flow field affects particles.

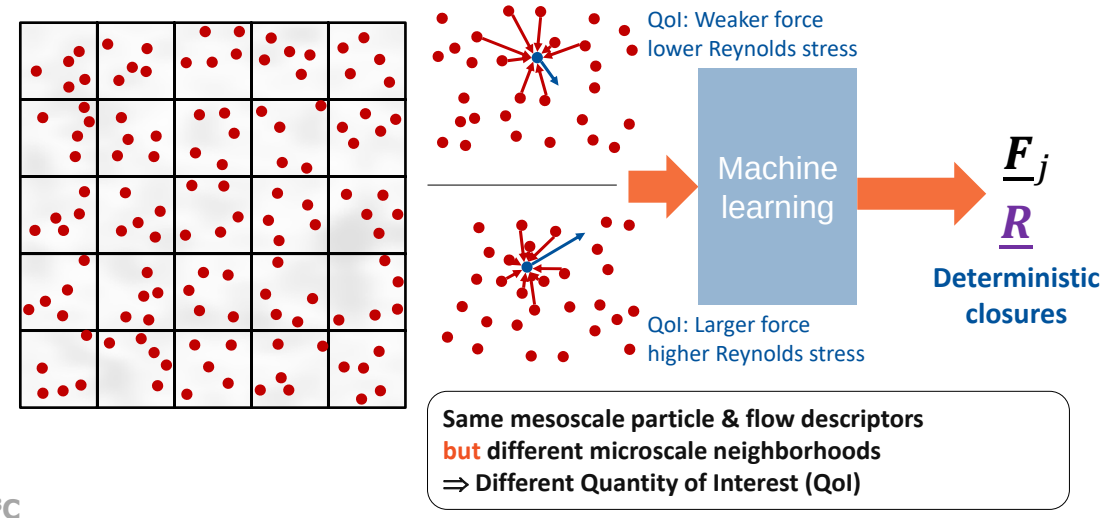

Same mesoscale particle & flow descriptors
but different microscale neighborhoods
⇒ Different Quantity of Interest (QoI)

Motivation for ML Models

- **Challenge:** Modelling every possible configuration explicitly is unreasonable

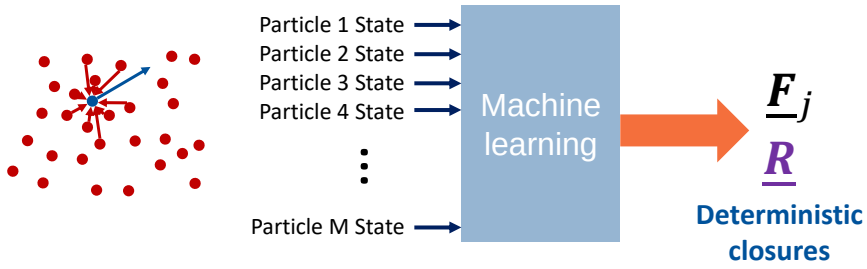

- **Goal:** Create a fast neural network that can account for <u>multi-particle dynamics</u>

Presentation Plan

- Motivation
- **Background**
- Overarching novel ML idea
- Data-Driven progress to date
- Physics-Informed modeling
- Next steps and other ideas

Prior Machine Learning Work

- **Direct force prediction (one step)**



- **How does the literature handle multi-particle dynamics?**
  - Want fewer data to train (only O(1000) PR-DNS particle forces)

Prior Machine Learning Work

- **Direct force prediction (one step)**

Configuration-specific Neighborhood info → (Binary / trinary / higher order effects) → Quantity of Interest

- **Superposition (Neural Operators):**
  - Siddani & Balachandar (PRF 2023, AICHE 2024)
    A. Wach's et al. (JFM 2020, IJMF 2025)
  - *Pairwise & trinary assumptions*
  - *Does not directly consider particle / fluid interactions*

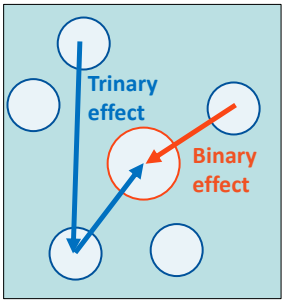

Prior Machine Learning Work

- **Direct force prediction (one step)**

Configuration-specific Neighborhood info → (Local Geometry) → Quantity of Interest

- **Local geometry (Convolutional Neural Network):**
  - Tafti et al. (Powder Tech 2019, 2022)
  - *Sees configurations but discretized and spatially limited*

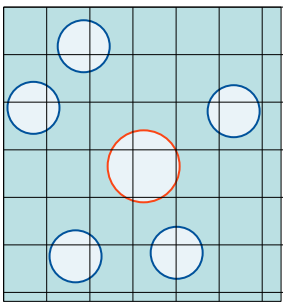

Prior Machine Learning Work

- **Direct force prediction (one step)**

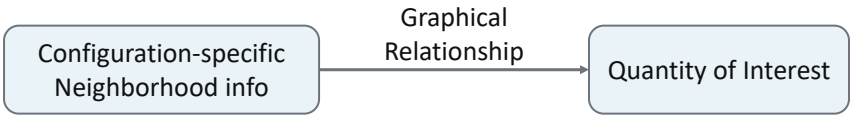

- **Relationships between particles (graph neural network):**
  - Ma et al. (Comp Meth Appl Math & Eng 2022, IJMF 2024)
  - *Establishes one-, two-, and three-body relationship*
  - *Allows for nonlinear relationships*
  - *Limited in spatial realizability*

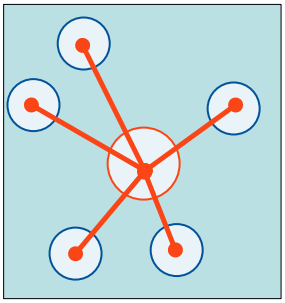

Flow Prediction as the Intermediary: Prior Work

- **Two-step force prediction**

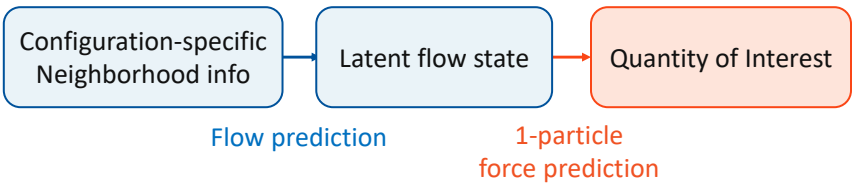

- **Learn to predict the PR-DNS flow**
  - Use this to predict the quantities of interest

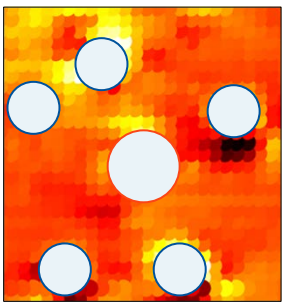

Flow Prediction as the Intermediary: Prior Work

- Akiki et al. (2017) — Pairwise Interaction Extended Point-Particle (PIEP) Model
  - Assumes linear interactions / disturbance fields
- Moore et al. (2020) — Data-Driven Superposable Wake
  - Needs to reconstruct flow field to estimate forces
- Stokesian Dynamics — Physics-Based Superposition
  - Simpler problem than we address here

Common philosophy: Predict the hydrodynamic interactions, evaluate particle forces second

Presentation Plan

- Motivation
- Background
- **Overarching novel ML idea**
- Data-Driven progress to date
- Physics-Informed modeling
- Next steps and other ideas

Key Innovation: Introduce Self-Consistent Learning

- **Two-step force prediction**

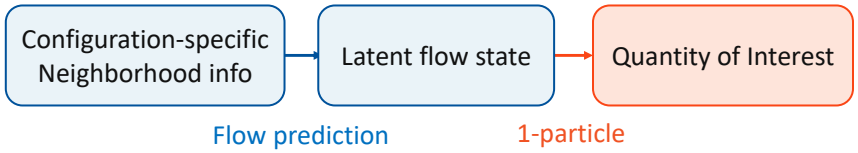

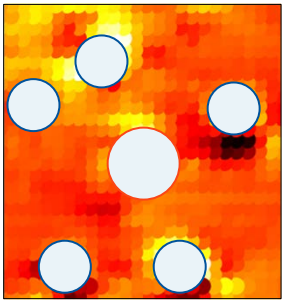

CM³C

Key Innovation: Introduce Self-Consistent Learning

- **Two-step force prediction**

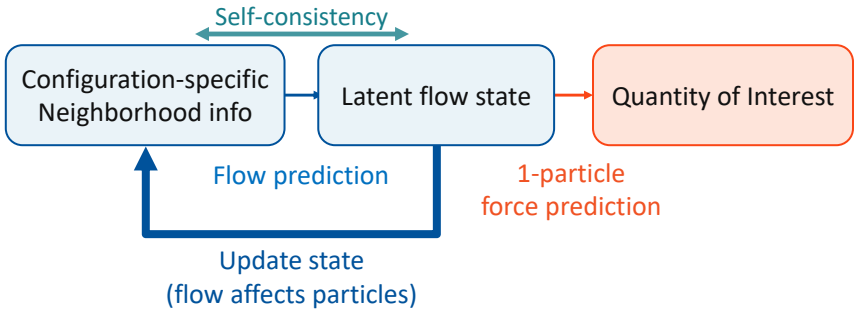

Superposition

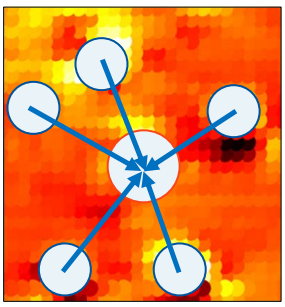

- **Refine the state**
  - Find the fixed-point solution

CM³C

UF UNIVERSITY of FLORIDA

Key Innovation: Introduce Self-Consistent Learning

- **Two-step force prediction**

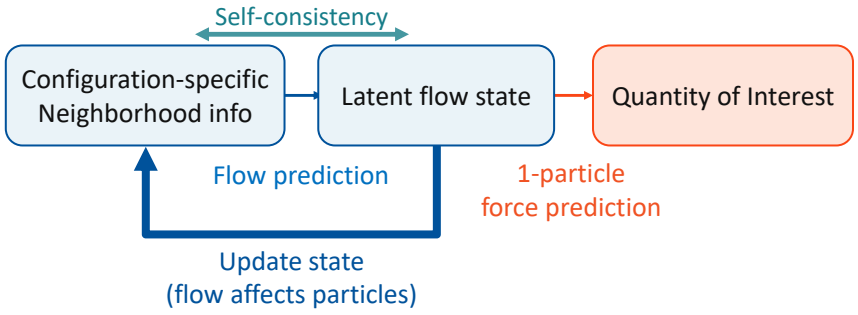

- **Refine the state**
  - Find the fixed-point solution

Superposition

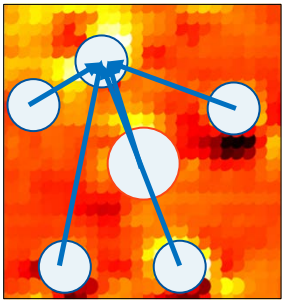

CM³C

UF UNIVERSITY of FLORIDA

Key Innovation: Introduce Self-Consistent Learning

- **Two-step force prediction**

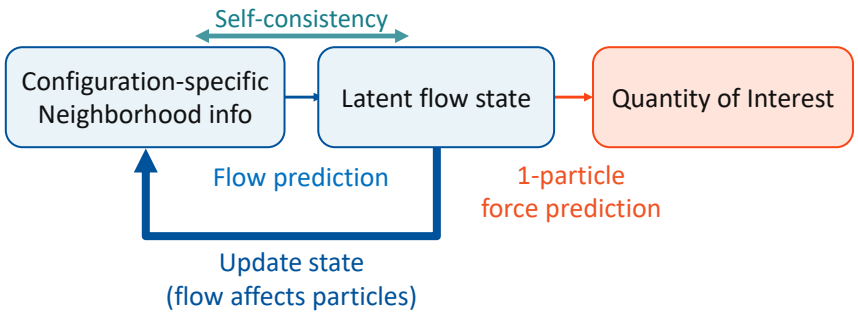

- **Refine the state**
  - Find the fixed-point solution

Refine Each Particle State

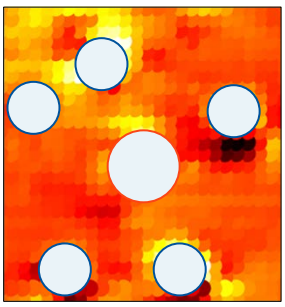

CM³C

UF UNIVERSITY of FLORIDA

Key Innovation: Introduce Self-Consistent Learning

- **Two-step force prediction**

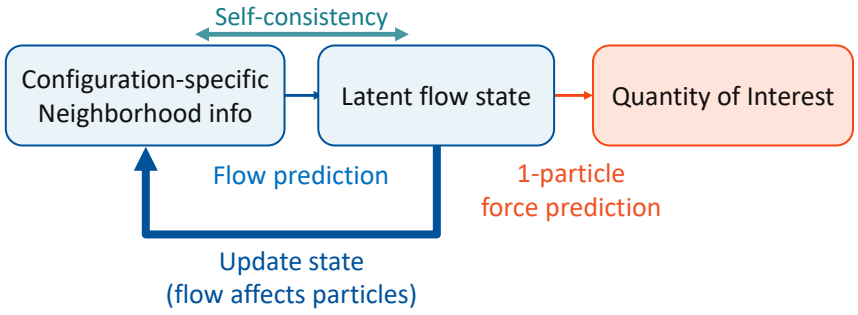

Superposition

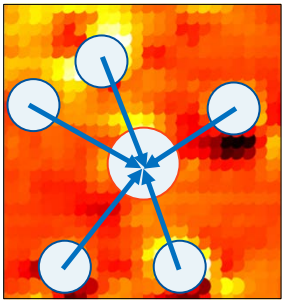

- **Refine the state**
  - Find the fixed-point solution

CM³C

UF UNIVERSITY of FLORIDA

Key Innovation: Introduce Self-Consistent Learning

- **Two-step force prediction**

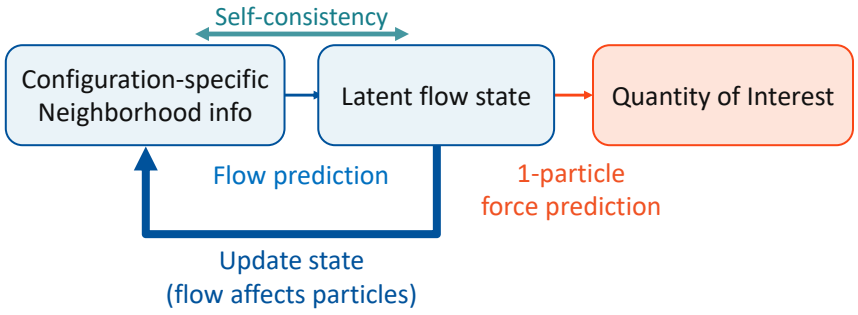

Superposition

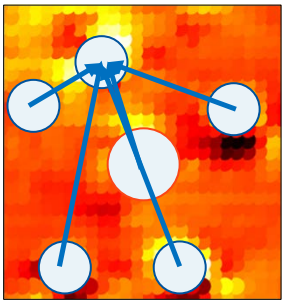

- **Refine the state**
  - Find the fixed-point solution

CM³C

Prior Evidence with Potential Flow

- **Flow Prediction-Based Deterministic Model**
  - Potential/linear flow problem where exact solutions are known
  - Neighbor-induced flow as a multipole expansion
  - Refining significantly improves the model

A cloud of evaporating droplets

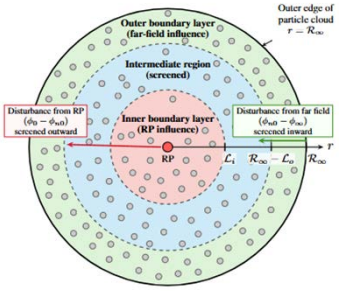

Average Model

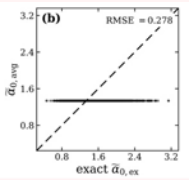

Stochastic Model

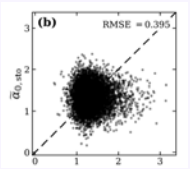

Iteratively Refined Model

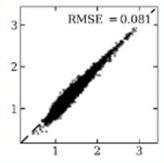

K=5 Neighbors

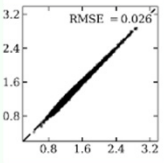

K=20 Neighbors

CM³C

Davin, A., Jacson, T., Balachandar, S., *Flow Prediction in Hydrodynamically Coupled N-Particle Systems via Superposition of One-Particle Solutions with Potential Flow Example*, Submitted JCP, 2026

Prior Evidence with Potential Flow

See Anderson Devin's poster for more information!

A cloud of evaporating droplets

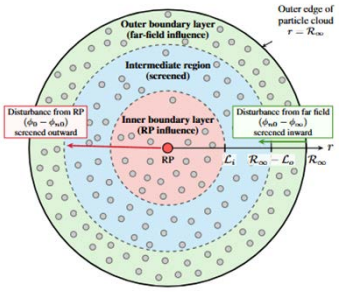

Average Model

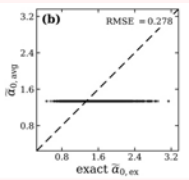

Stochastic Model

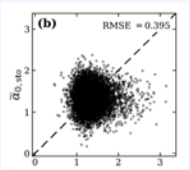

Iteratively Refined Model

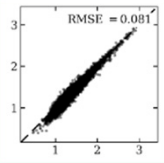

K=5 Neighbors

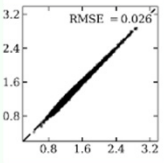

K=20 Neighbors

CM³C

Davin, A., Jacson, T., Balachandar, S., *Flow Prediction in Hydrodynamically Coupled N-Particle Systems via Superposition of One-Particle Solutions with Potential Flow Example*, Submitted JCP, 2026

Our Neural Network Model

- **Train Neural Network (neural operator)**
  - Minimize squared error (no physics constraints)

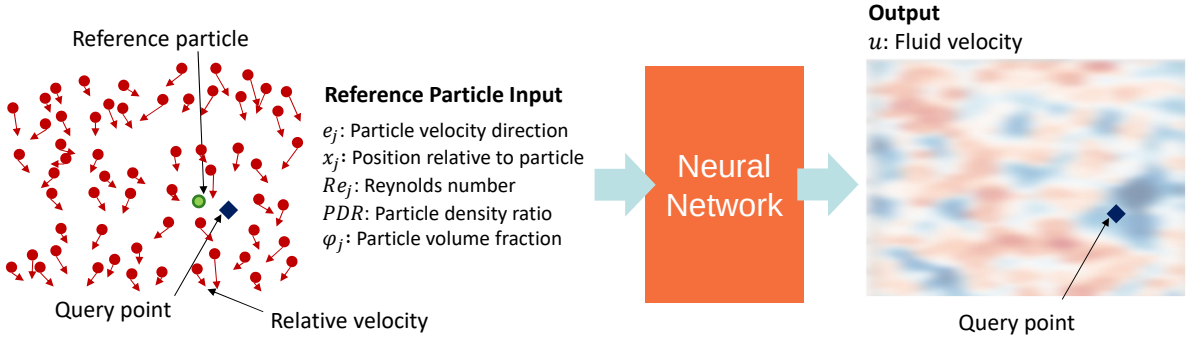

Our Neural Network Model

- **Train Neural Network (neural operator)**
  - Minimize squared error (no physics constraints)

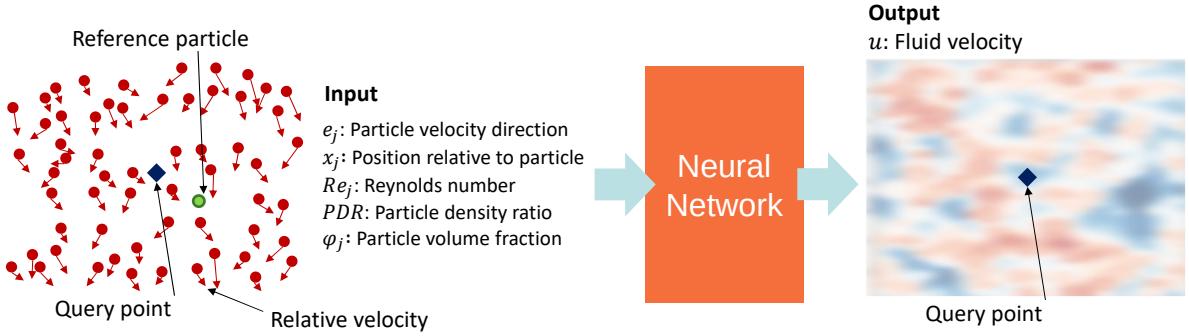

Repeat for many combinations

Our Neural Network Model

- **Train Neural Network (neural operator)**
  - Minimize squared error (no physics constraints)

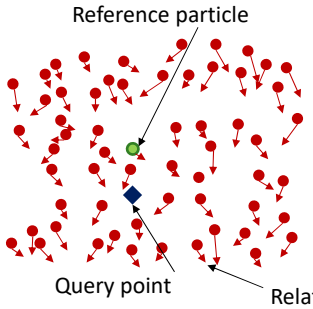

Reference Particle Input

e_j: Particle velocity direction
x_j: Position relative to particle
Re_j: Reynolds number
PDR: Particle density ratio
φ_j: Particle volume fraction

Neural Network

Output
u: Fluid velocity

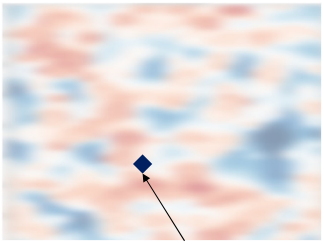

Repeat for many combinations

Our Neural Network Model

- **Learned model:** single-particle for given fluid / particle state
  - From a diverse collection of examples

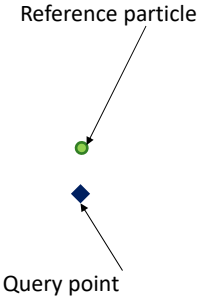

Reference Particle Input

e_j: Particle velocity direction
x_j: Position relative to particle
Re_j: Reynolds number
PDR: Particle density ratio
φ_j: Particle volume fraction

Neural Network

Output
u: Fluid velocity

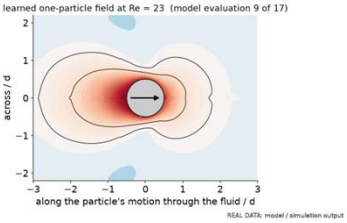

Raster query points

Our Neural Network Model

- **Use superposition to approximate the field**
  - It will have errors since flow is not a superposition

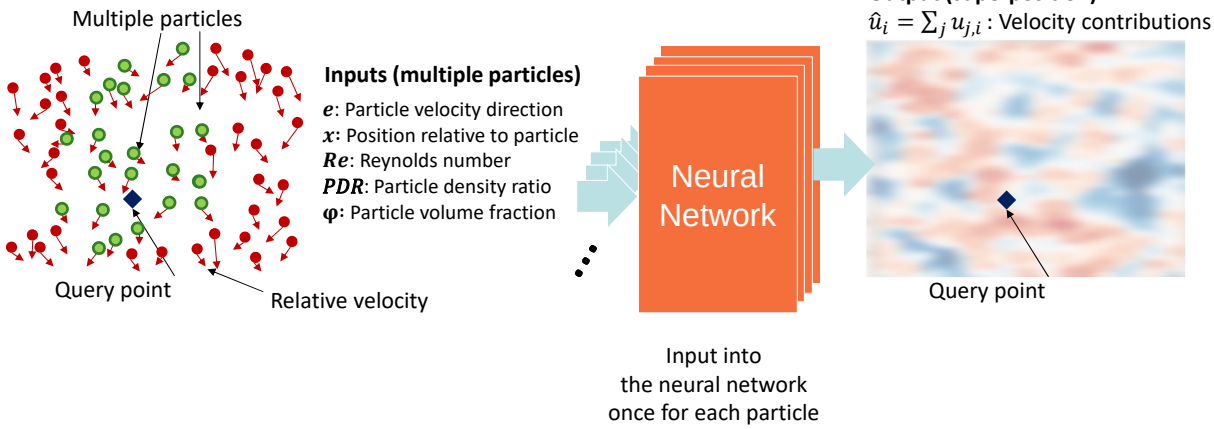

Our Neural Network Model

- **Update Particle**
  - Refine particle's relative velocity

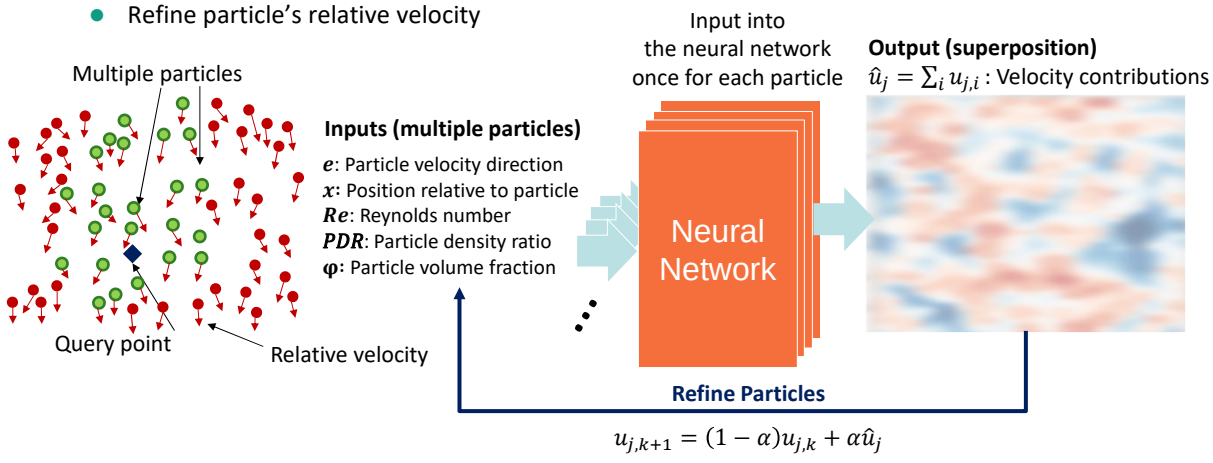

Our Neural Network Model

- **Update the flow**
  - With new particle information

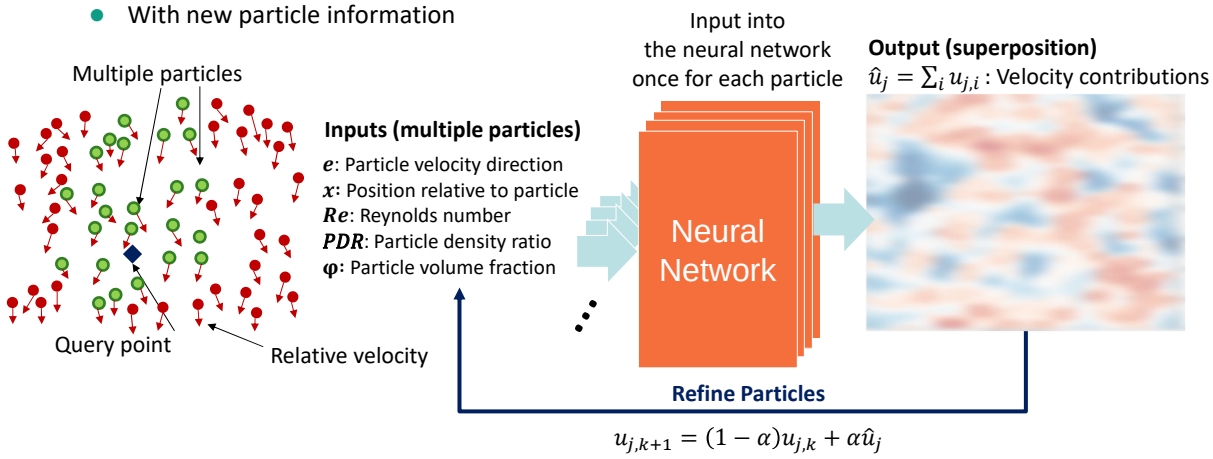

Our Neural Network Model

- **Iterate** until toward a fixed point (equilibrium learning in ML literature)

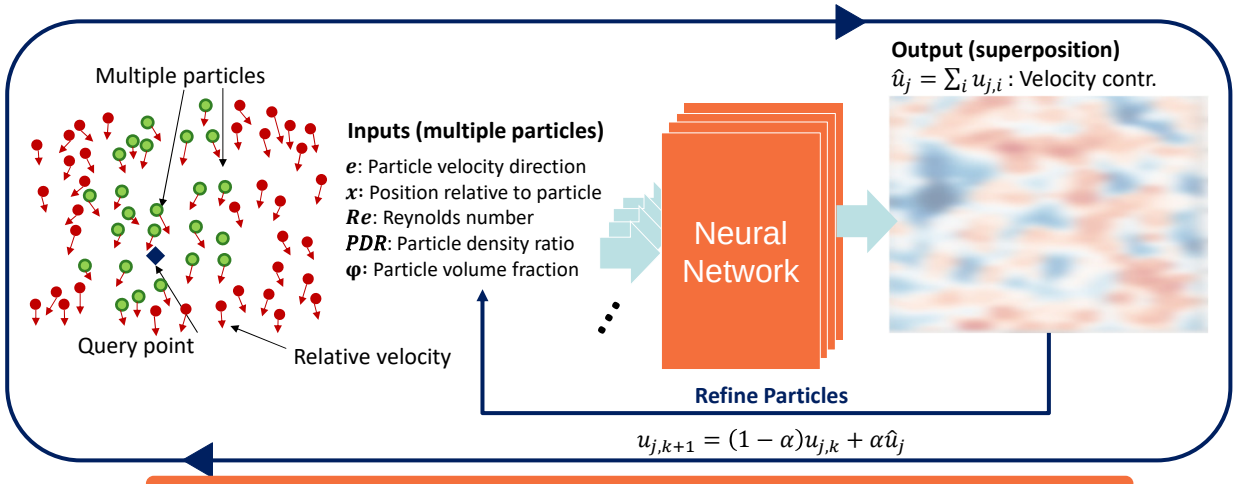

Achieves self-consistency between the particles

Our Neural Network Model

- **Learned model:** single-particle <u>expansion</u> for given fluid / particle state

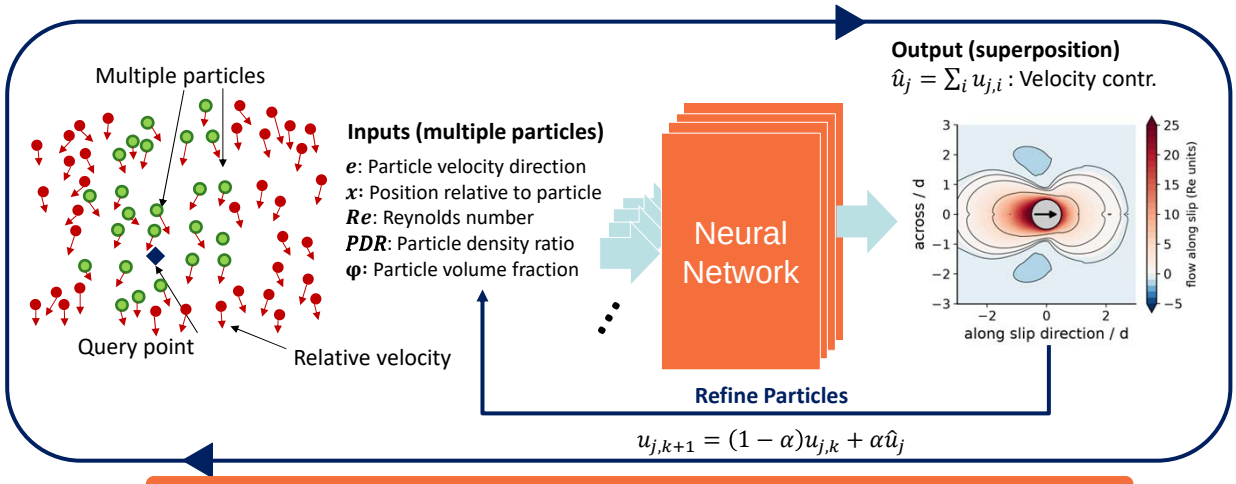

Achieves self-consistency between the particles

Our Neural Network Model

- **Also apply this during inference** (not common but very important)

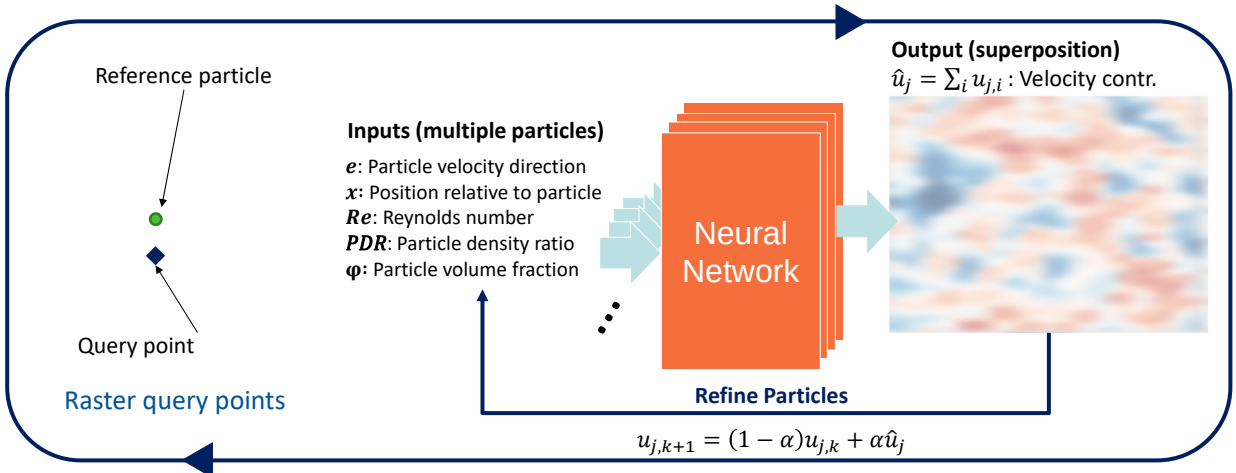

Achieves self-consistency between the particles

Presentation Plan

- Motivation
- Background
- Overarching novel ML idea
- **Data-Driven progress to date**
- Physics-Informed modeling
- Next steps and other ideas

Simulated Data

- **Particle-Resolved Simulation (General Info)**
  - Particles falling in otherwise stationary fluid
  - Triply periodic boundaries
  - Rigid bodies, no rotation
  - Data are non-dimensionalized
    - Length scale: d (particle diameter)
    - Velocity scale: d / v (diameter / kinematic viscosity)
  - Mesh dimensions: $150 \times 300 \times 150$
    - $10 \times 20 \times 10$ particle diameters
  - Variety of Re, PDR, and average φ

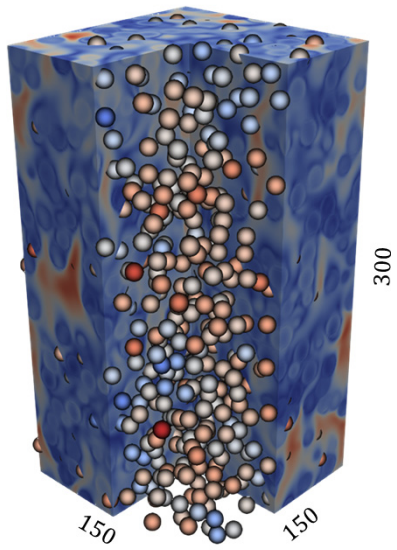

Simulated Data

- **ML-Used Data**
  - One simulation (one Re, PDR, and average φ)
  - 64 particles
  - Training with 9 snapshots
  - Validation / hyperparameter tuning with 3 snapshots
  - Testing with 3 snapshots

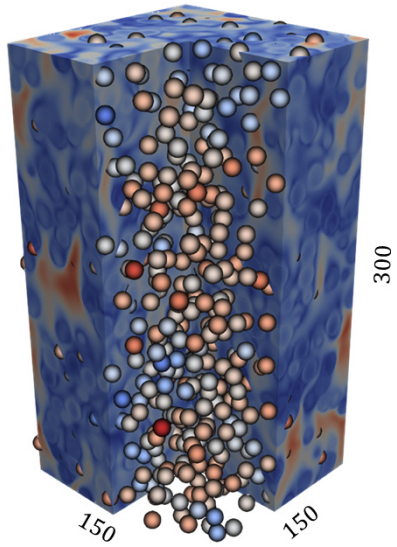

Understanding the Data Before ML Modeling

- **Case 13:**
  - $\langle \mathrm{Re} \rangle = 10$
  - $\rho = 2.56$
  - $\langle \phi_p \rangle = 26.22\%$
  - Large drag tail
  - Strong negative between drag force and shear force
  - Force clusters

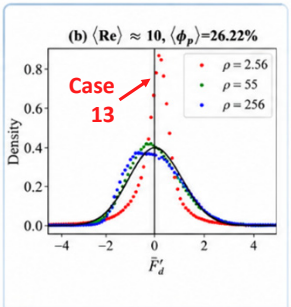

Drag PDF at high $\langle \phi_p \rangle$
Case 13 develops a pronounced negative tail.

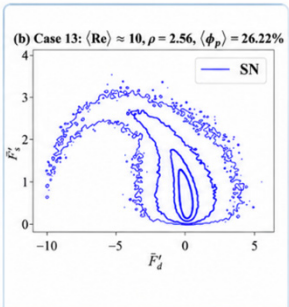

Case 13 jPDF
Strong negative correlation: $\rho_{sd} \approx -0.50$.

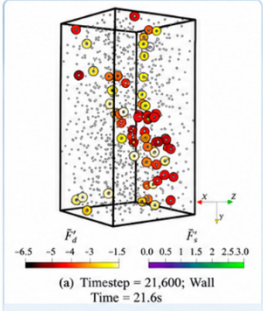

Transient structure
Dynamic anisotropic clusters linked to extreme force events.

Lott, C., Zhang, D., Wang, M., Balachandar, S., *JFM (under review)*, 2026.

Machine Learning Architecture

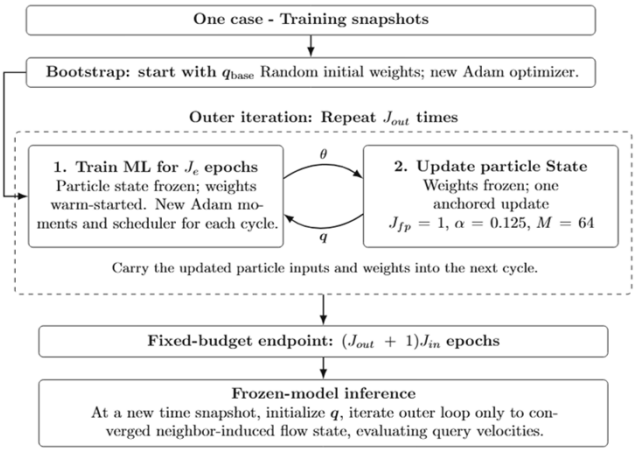

Machine Learning Architecture

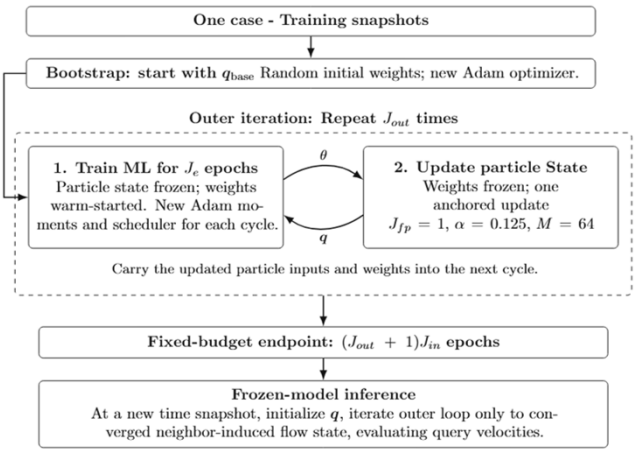

- **Neural Network Architecture**

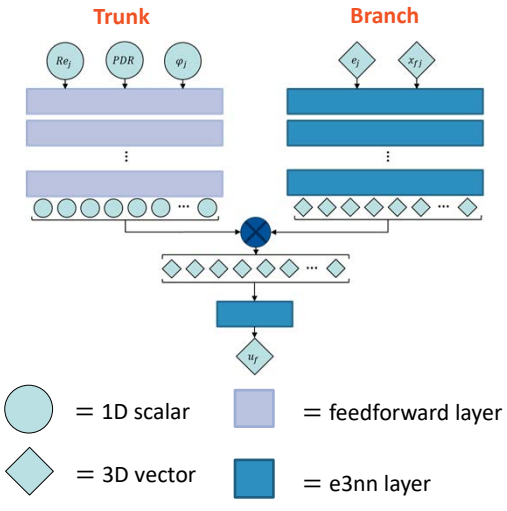

Hyperparameter Tuning

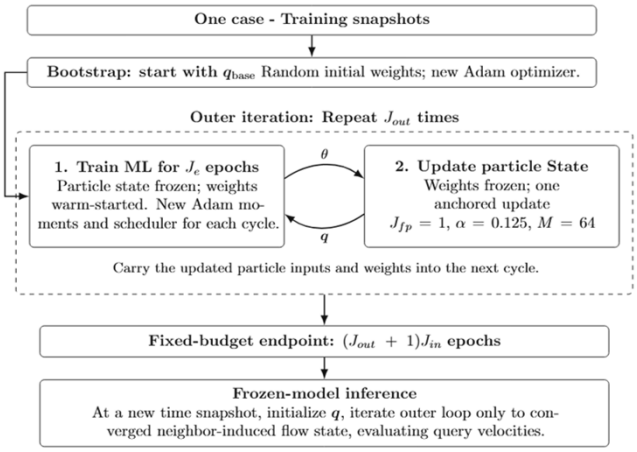

- **Hyperparameters Searched**
  - Case 13, 32 seeds, 1,152 runs
  - Training epochs per cycle J ∈ {1, 2, 4};
  - State updates per cycle K ∈ {1, 2, 4};
  - Damping α ∈ {1/8, 1/4, 1/2, 1}.
- Selected
  - J = 4, K = 1, α = 1/8

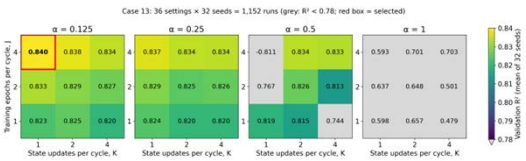

  - **Validation:** R^2 0.840 ± 0.003). Damping matters most: α ≥ 1/2 is unstable

CM³C

What does the reconstructed flow look like?

- **A held-back test Case 13**

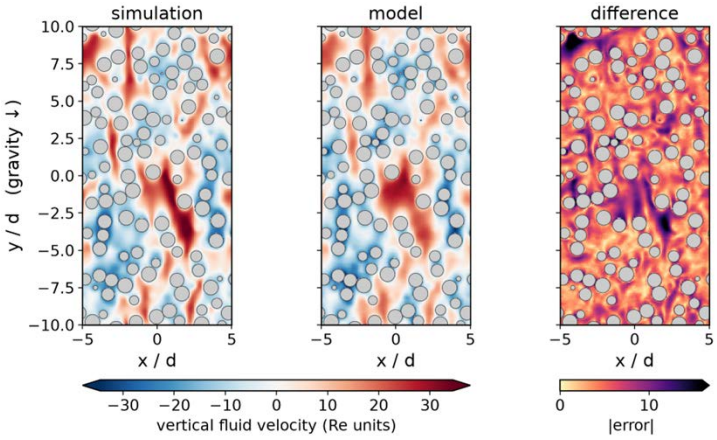

  - **Self-Consistent slice accuracy:** R^2 = 0.83
  - **Self-Consistent whole box accuracy:** R^2 = 0.84 (vs 0.74 for the fixed-state model)

CM³C

What does the reconstructed flow look like?

- **A held-back test Case 13**

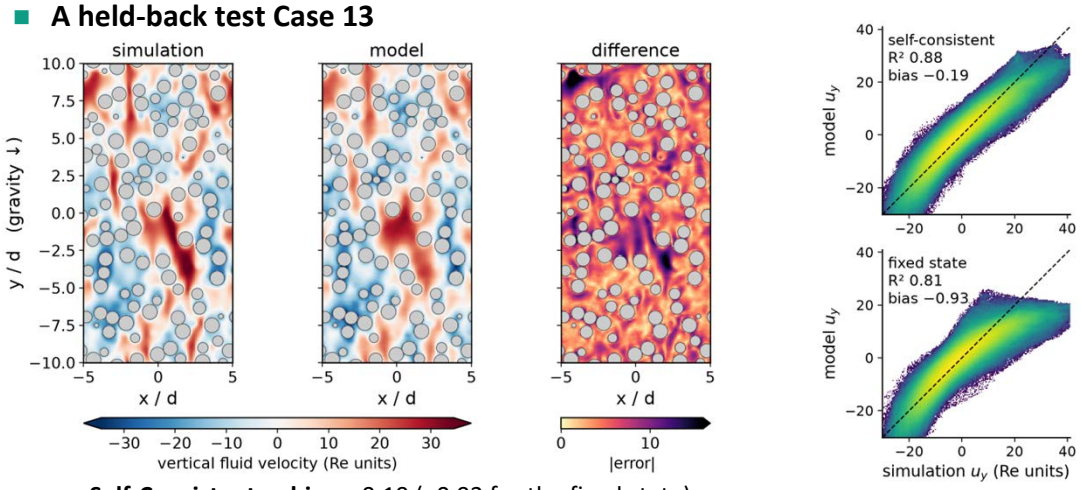

- **Self-Consistent u_y bias:** −0.19 (−0.93 for the fixed state)
  - **Under 0.2 in every component.** The largest misses are at extreme velocities.

CM³C

Refining the initial flow estimate

- **Self-consistency improves performance on snapshots the model never saw**

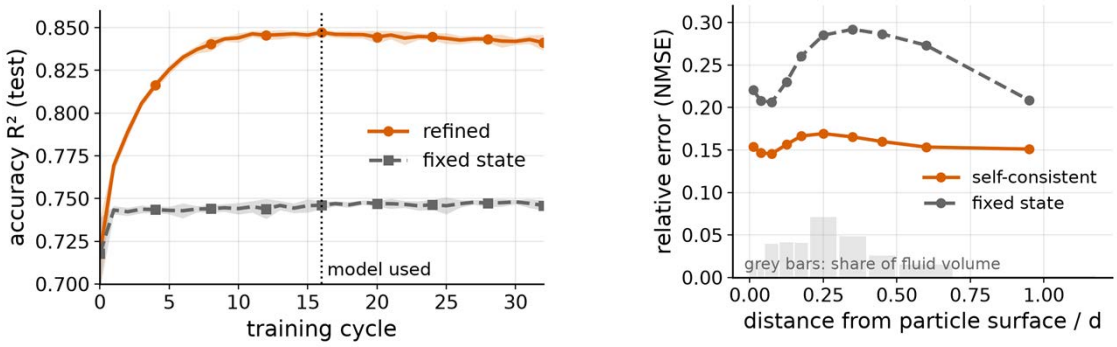

CM³C

How many refinement steps when predicting?

- **Freeze the trained model and refine from each particle's own velocity**

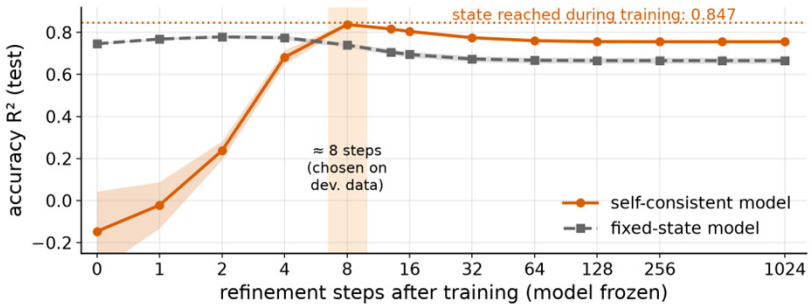

- Without refinement, the self-consistent model is poor (R^2 −0.15).
- Full convergence settles lower (0.76).

Where the gain comes from

See Raul Valle's poster & talk for more information + results!

What about the uncertainty of our prediction?

- **A held-back test Case 13**

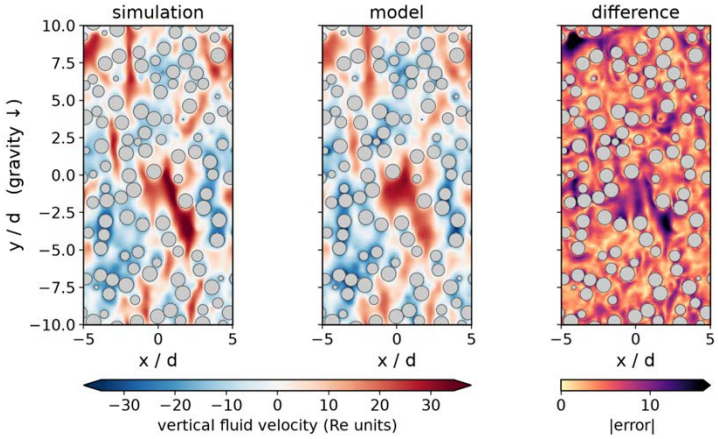

What about the uncertainty of our prediction?

- **A held-back test Case 13**

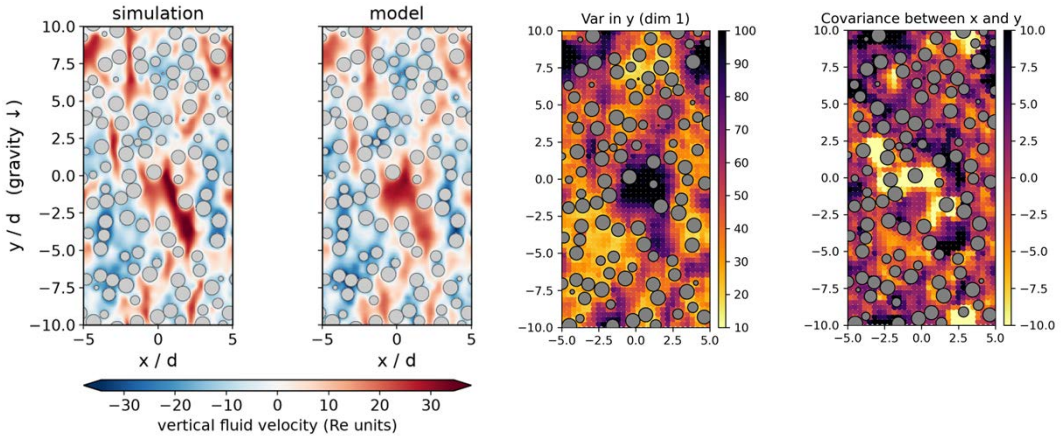

- **Initial measure of <u>epistemic uncertainty</u> (fixed state only)**
  - Highest in <u>the fluid between particles</u> where neighbors' flows overlap

What does the reconstructed flow look like?

See Isaly Tappan's poster for more information + results!

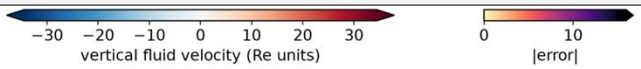

- **Initial measure of <u>epistemic uncertainty</u> (fixed state only)**
  - Highest in the fluid between particles where neighbors' flows overlap

Presentation Plan

- Motivation
- Background
- Overarching novel ML idea
- Data-Driven progress to date
- **Physics-Informed modeling**
- Next steps and other ideas

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

T4: Physics-Informed Machine Learning

Juan Diaz and Alina Zare

Department of Electrical and Computer Engineering

University of Florida

UF UNIVERSITY of FLORIDA PURDUE UNIVERSITY The University of Texas at Austin Oden Institute for Computational Engineering and Sciences

UF UNIVERSITY of FLORIDA

This work

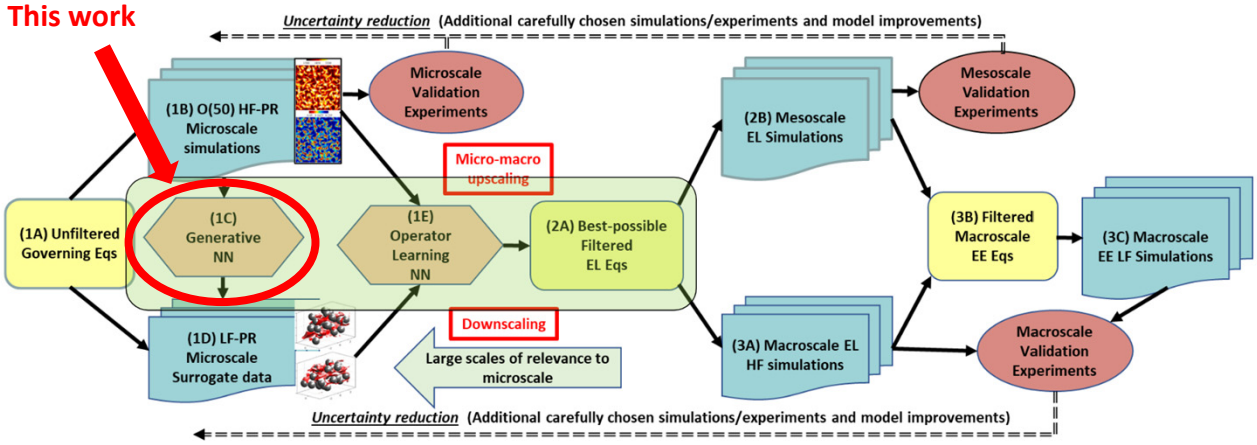

Microscale Mesoscale Macroscale

Problem Description

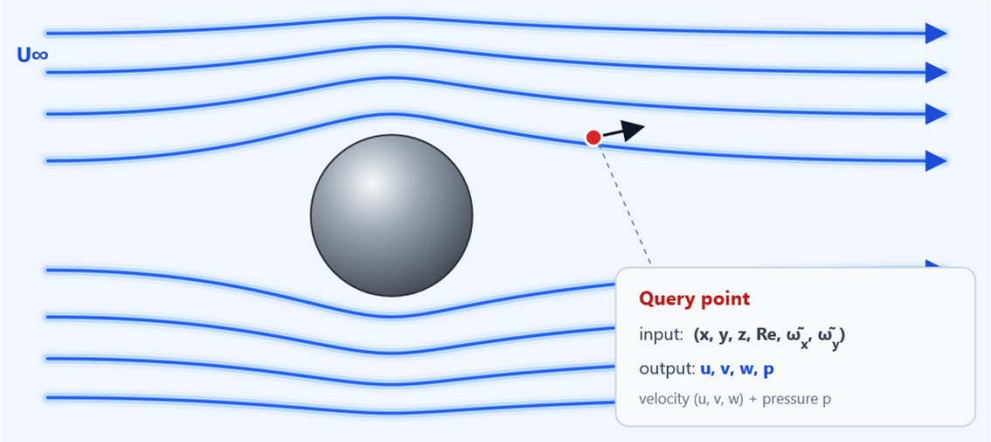

Setting: Sphere in a fluid
Inputs: Location, Reynold's number, Rotation in x- and y-
Outputs: Velocity and Pressure

Physics-Informed Neural Network (PINN)

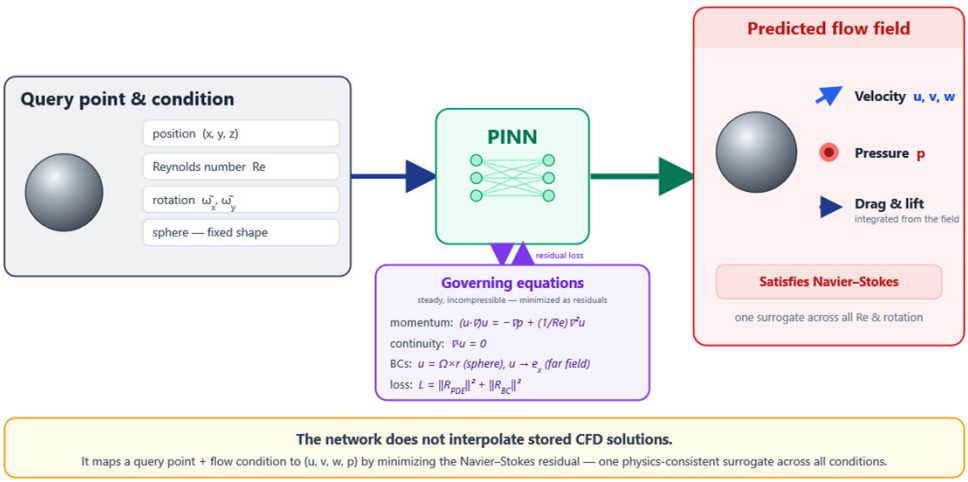

Standard PINN: MLP

Single Case: Re = 50 , rotation (0,0)

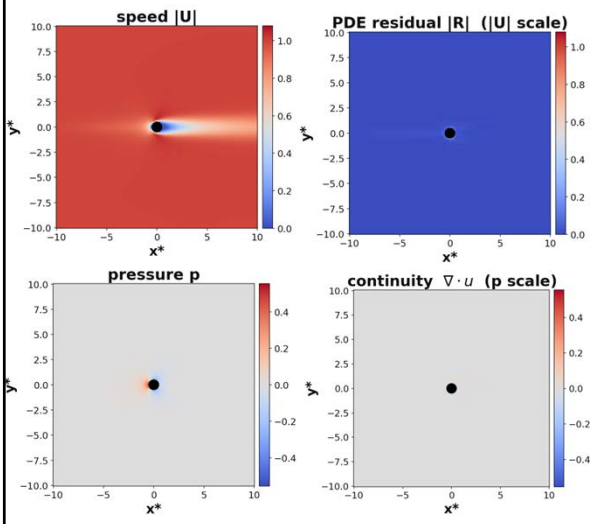

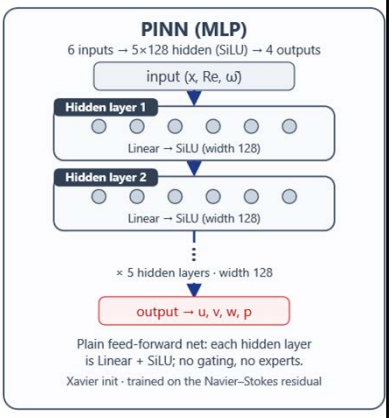

CM³C

Standard PINN: MLP

Single Case: Re = 50 , rotation (0,0)

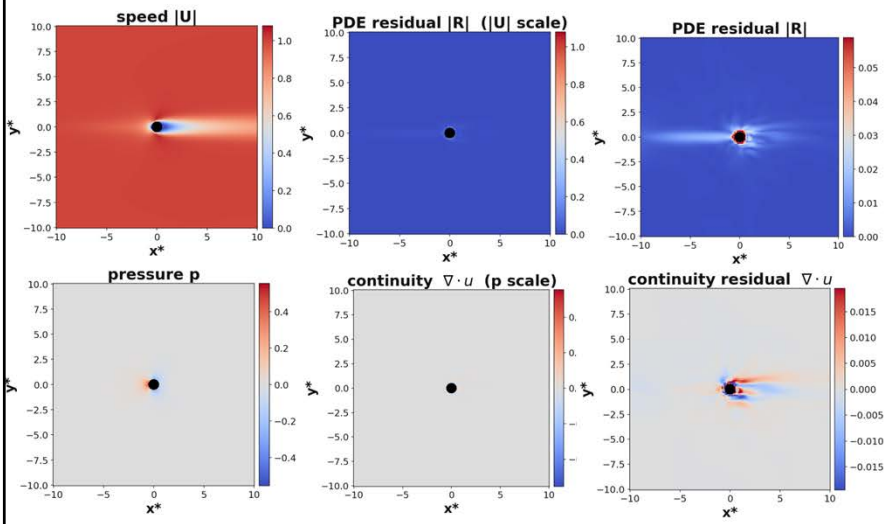

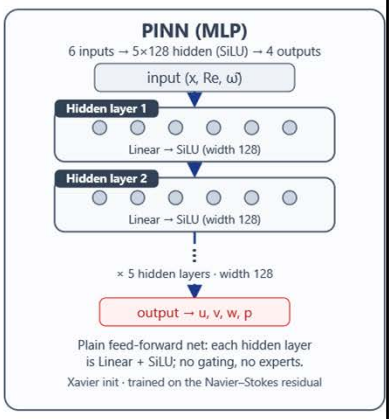

CM³C

Standard PINN: Performance Across Re

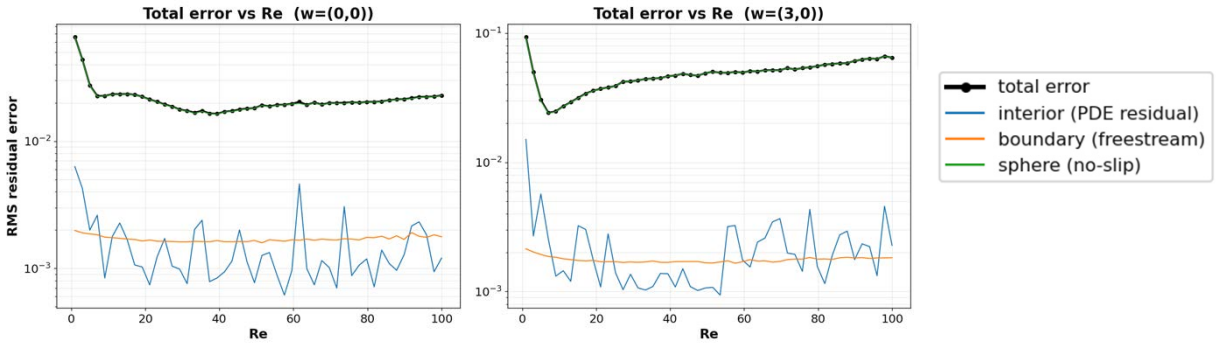

Mixture-of-Experts (MoE): 35 Years of Research

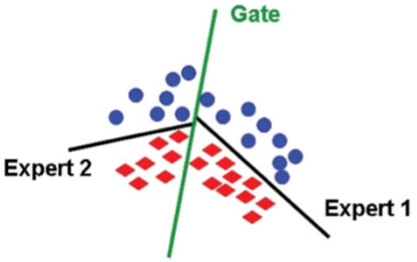

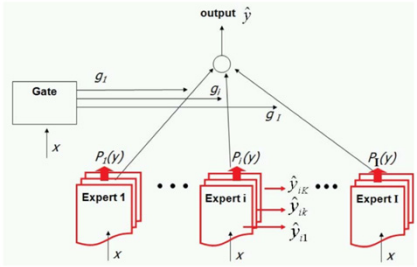

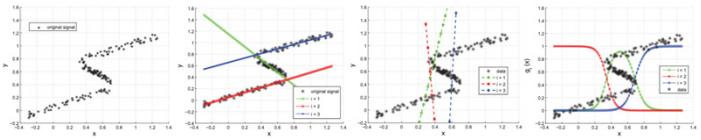

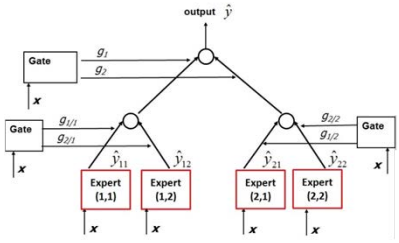

Figures from: Twenty Years of Mixture of Experts, 2012, Seniha Esen Yuksel, Joseph N. Wilson, and Paul D. Gader

Dense MoE : Unguided Soft Gating Mechanisms

Single Case: Re = 50 , rotation (0,0)

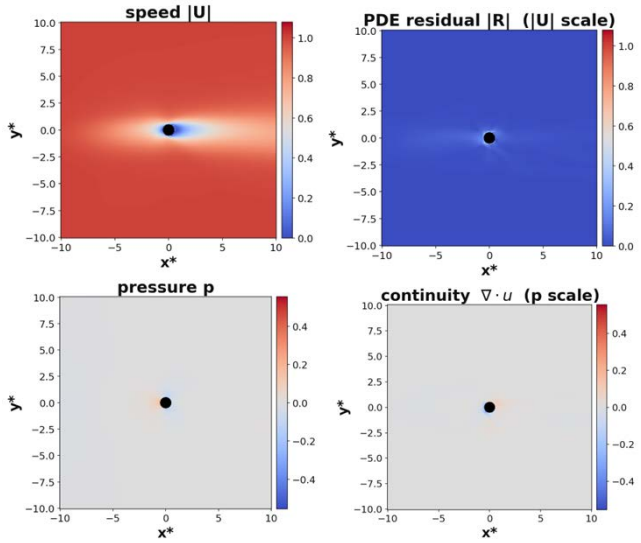

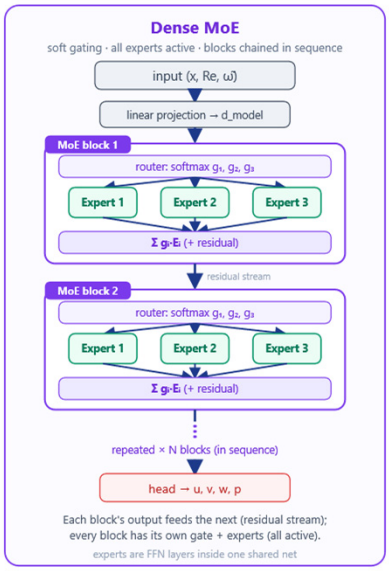

CM³C

Distance MoE: Hard Gating Constrained by Distance

Single Case: Re = 50 , rotation (0,0)

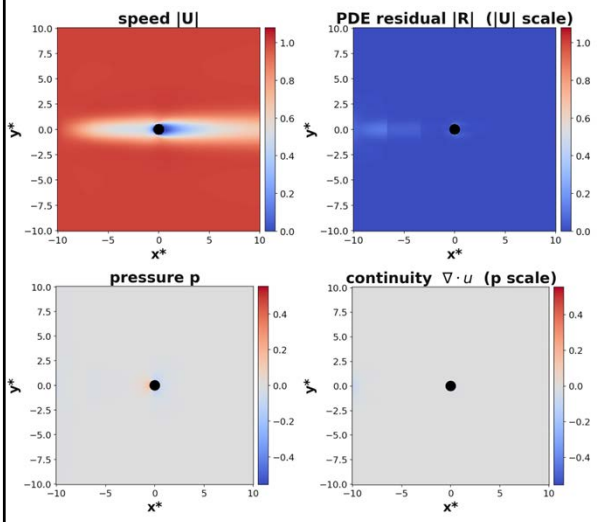

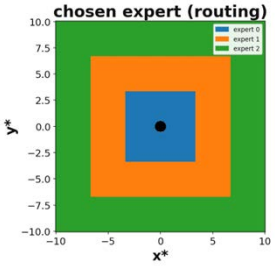

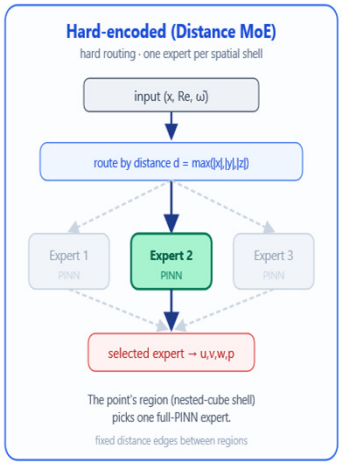

CM³C

Re MoE: Band Guided Gating Mechanisms

Single Case: Re = 50 , rotation (0,0)

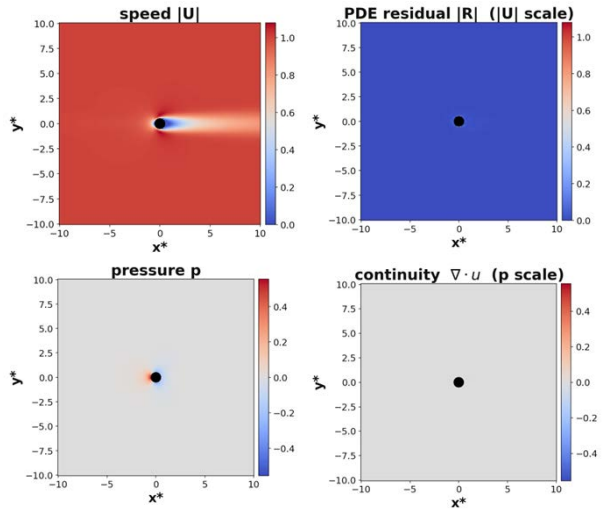

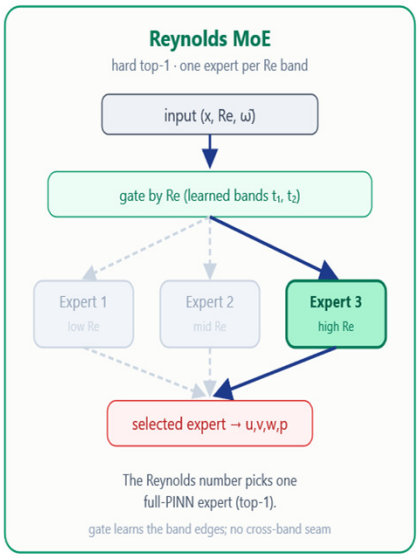

Re MoE: Band Guided Gating Mechanisms

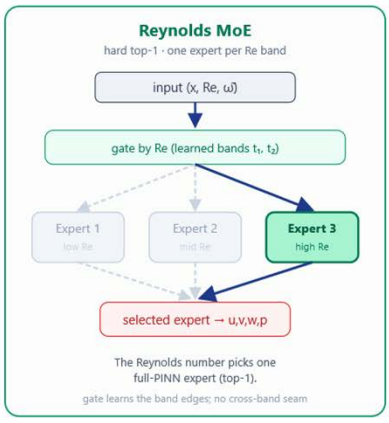

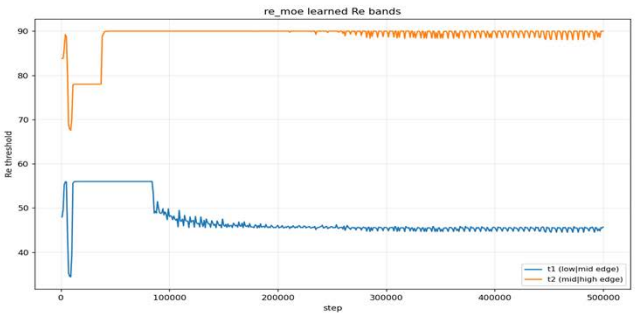

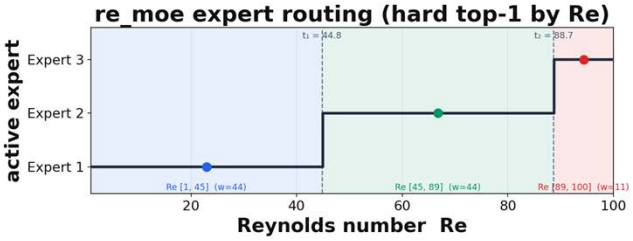

PINN Comparisons

No rotation
sweep across Re
range [1,100]

MLP (blue) is baseline

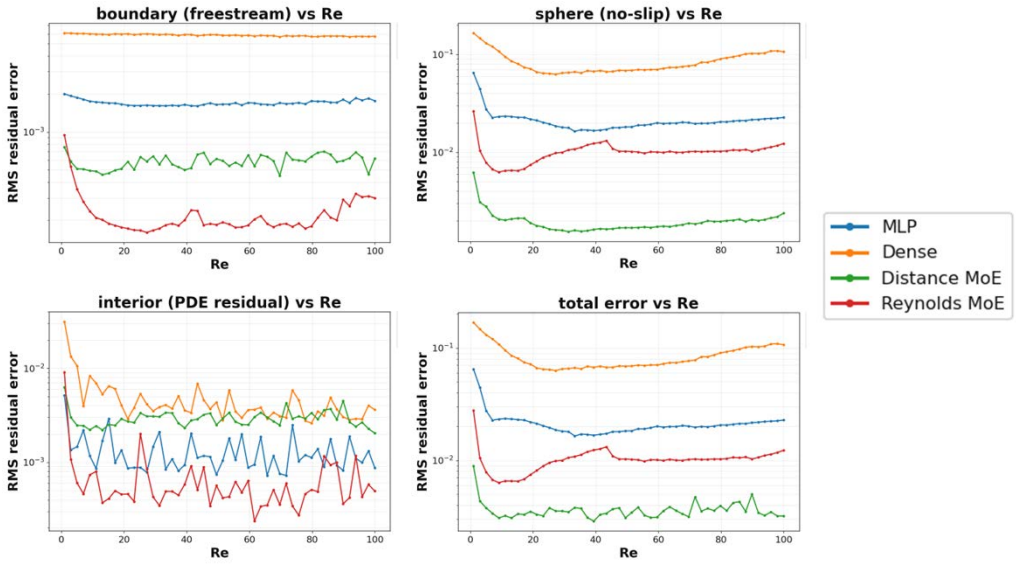

Expanding PINN Capabilities

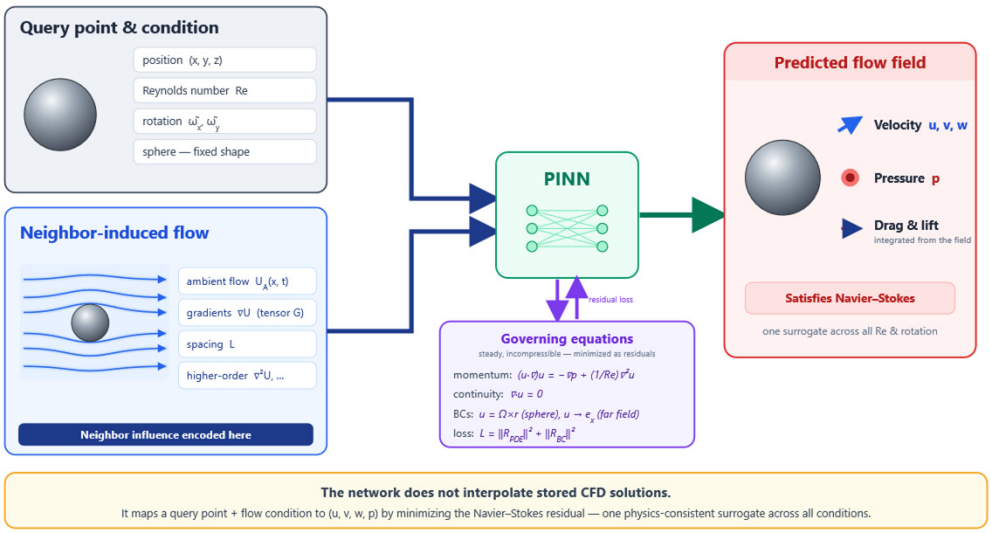

Neighbor Induced Flow PINN: MLP

Re = 50, rotating sphere (ω̃_y = 0.5) in a neighbor-induced shear field (G_yx = 0.25) at close spacing (L = 1.5)

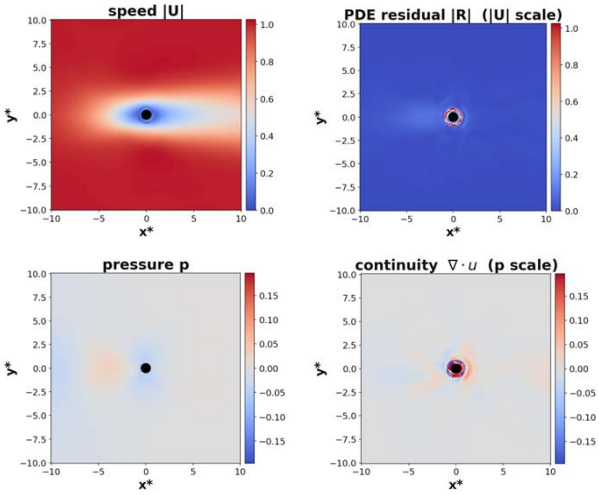

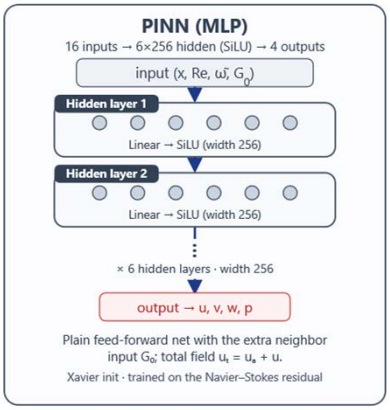

Neighbor Induced Flow PINN: Distance MoE

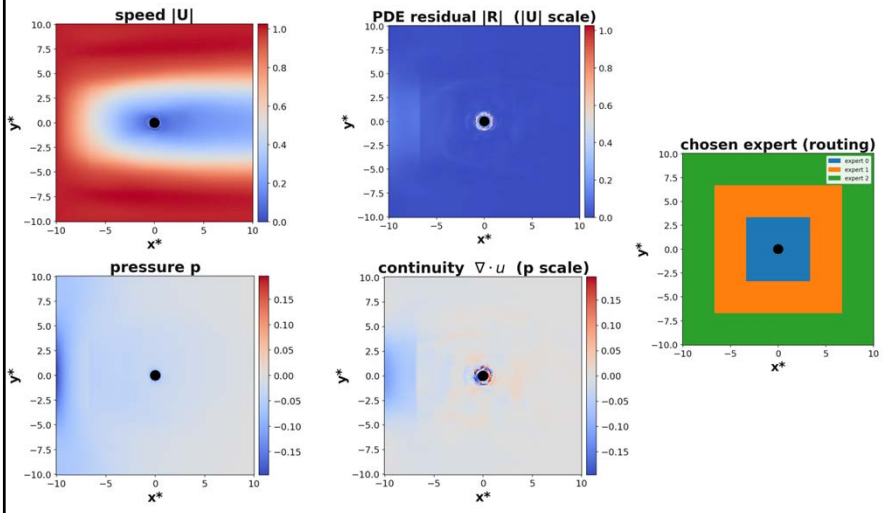

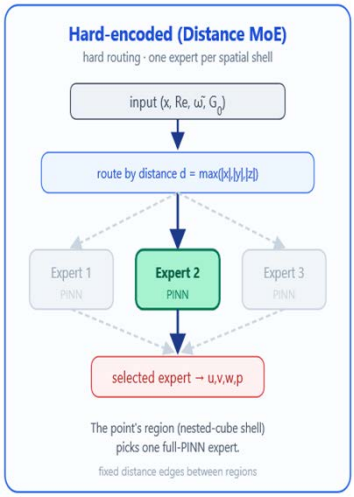

Neighbor Induced Flow PINN: Reynolds MoE

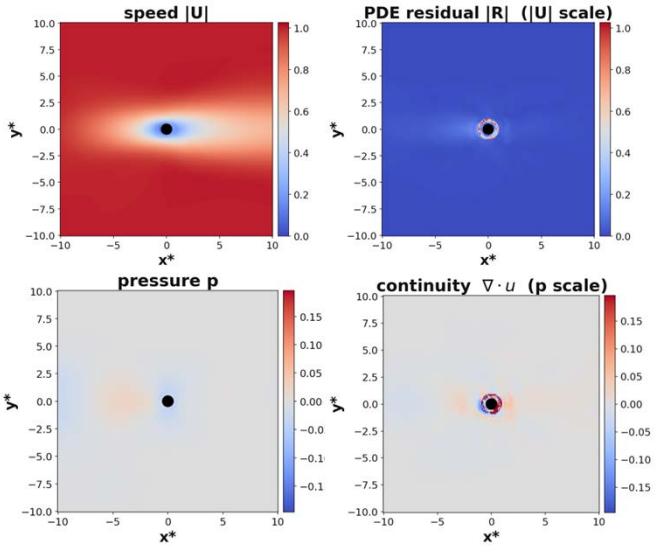

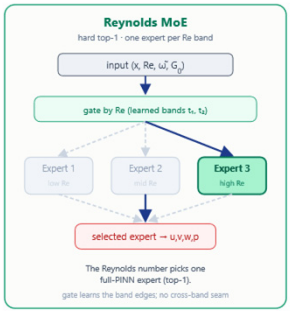

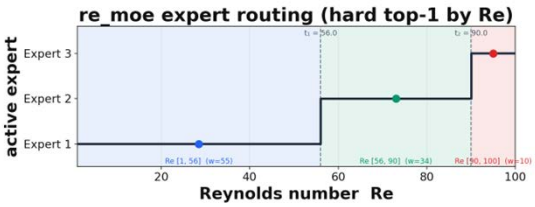

CM³C

Neighbor Induced Flow PINNs Across Re

Rotating sphere (ω̃_y = 0.5) in a neighbor-induced shear field (G_yx = 0.25) at close spacing (L = 1.5)
Swept across Re range [1,100]

MLP (blue) is baseline

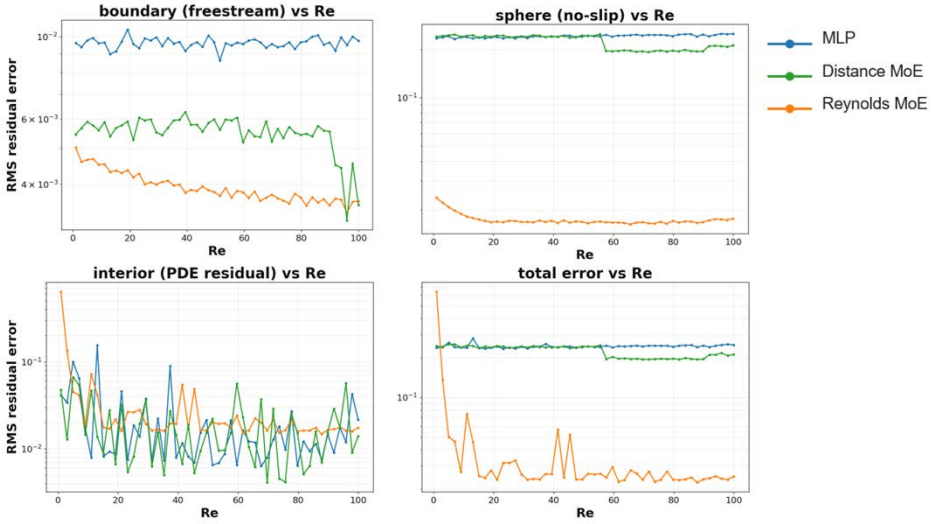

CM³C

Summary

- PINN built from existing literature
- Expanded by:
  - Applying Mixture of Experts
  - Introducing neighbor induced flow
- MoE PINNs with guided gates outperform standard PINN

Future Work

- Optimize MoE:
  - Soft Gating
  - Improved Expert Networks
  - Multiple Gating Guides
  - Improved Sampling per expert

Introduction Machine Learning & Sensing Lab





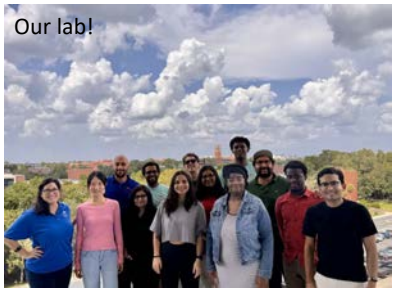



CM³C

Thank you.

Questions?



NNSA
National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

T5. Experiments and V&V-UQ, Part 1

Co-Principal Investigator: D. Guildenbecher

School of Mechanical Engineering

Purdue University

UF UNIVERSITY of FLORIDA PURDUE UNIVERSITY The University of Texas at Austin Oden Institute for Computational Engineering and Sciences

PURDUE UNIVERSITY

CM³C Experimental Plans

Demonstration Scale ⟷ **Mesoscale** ⟷ **Microscale**

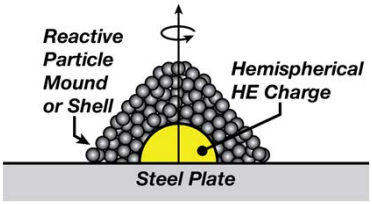

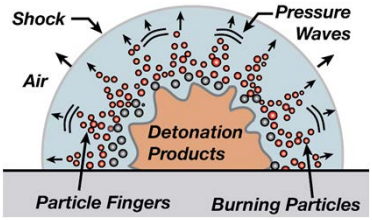

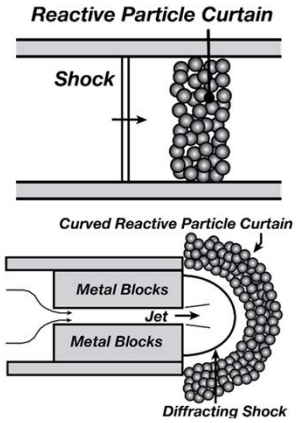

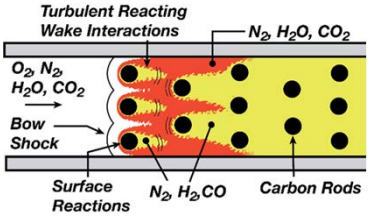

CM³C Experimental Schedule

Year 1-2 Tasks:

- *Establish the demonstration problem* → Explosive tests to be performed at Purdue using our existing facilities and diagnostics
- *Lay the foundation for HF-PR validation* → Design/build carbon rod experiments and validate single rod results against published literature

Task	Year 1	Year 2	Year 3	Year 4	Year 5
Explosive demonstration using standard imaging and pressure gauges			▲		
Single rod experiments validated against published literature		▲			
Design/build upgraded multi-rod configurations					
Multi-rod group combustion measurements				▲	
Sub-scale testing of blast-hardened LAS/emission/extinction gauges					
Design/build multi-rod configurations with spacing guided by UQ team					
Multi-rod, multi-parameter experimental measurements guided by UQ team					▲
Demonstration experiments with LAS and other advanced diagnostics				▲	
Final Publication/documentation					▲

▲ Indicate major milestones, see T1 presentation by S. Balachandar

CM³C

CM³C Experimental Team



Cohen Nunes, 3rd year PhD student

- Expected graduation Spring or Summer 2027
- Leading optical diagnostics for demonstration scale experiments
- Supported by PSAAP



Bekah Travis, 4th year DPhD student

- Expected graduation Spring or Summer 2027
- Leading explosive physics for demonstration scale experiments
- NDSEG funded and SNL graduate student intern



Alex Millauer-Hsu, 1st year DPhD student

- Expected graduation ~2030
- Leading microscale c-rod experiments
- Supported by PSAAP



Colton Willhardt, Postdoc

- Supporting experimental operations
- Prior SNL graduate student intern
- Partially supported by PSAAP (~15%)

Assisted by several others who are supported on related projects: Nick Gurdian-Aponte (UGrad research for credit), Rees Verleur (3rd year DPhD, supported by NNSA/SSAP), Amogh Natu (2nd year DPhD, supported by AFRL), Amy Chang (2nd year PhD, supported by NASA), Daniel Struk (1st year DPhD supported by fellowship and ONR)

CM³C

CM³C Demonstration Problem

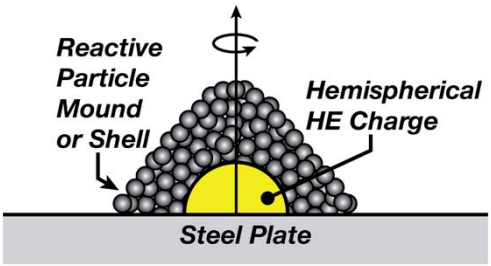

Charge: 3 cm diameter hemisphere of PETN ($C_5H_8N_4O_{12}$)

Particles: Free-standing mounds of carbon and magnesium particles

- Carbon reacts at the surface. Magnesium reacts in a vapor flame

Diagnostics: ***QoI1*:** high-speed imaging for blast wave locations; ***QoI2*:** pencil gages for pressure; ***QoI3*:** visible and X-ray extinction for particle concentration

Advantages:

1. Well-defined initial conditions, free of uncertainty due to case confinement
2. PETN equations of state (EoS) are publicly available, and PETN produces little soot
3. We have extensive modeling and experimental experience with this base charge



Ongoing Metalized Charge Experiments At Purdue (Courtesy S. Son, R. Travis)

CM³C Demonstration Problem

Our diagnostic advancements will enable measurements of ***QoI4*:** temperature variations due to combustion and heat release, ***QoI5*:** fingering instabilities, and ***QoI6***: PDFs of pressure and particulate volume fraction

Blast-Hardened Laser Absorption Spectroscopy

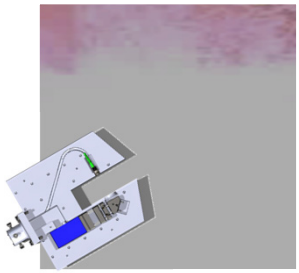

Travis *et al*. (2025) *Work In Progress*

Quantitative Schlieren

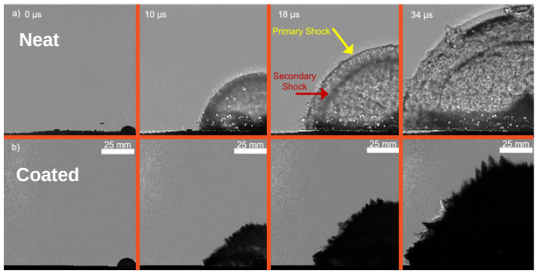

Nunes and Guildenbecher *et al.* (2025) *AIAA SciTech*

Extinction Imaging

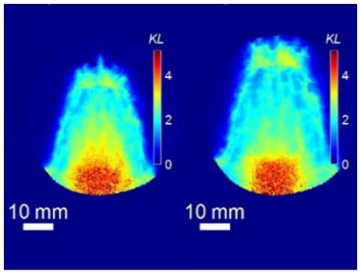

Saltzman and Guildenbecher *et al.* (2023) *Propellants Explos. Pyrotech.*

CM³C Demonstration Problem

Planned DoD partnerships create extensive opportunities for larger scale demonstration experiments



Naval Surface Warfare Center, Crane Division, Crane, Indiana

- Indoor and outdoor explosive test facilities focused on expeditionary warfare missions
- Extensive collaborations with Purdue

Air Force Research Laboratory, Munitions Directorate, Eglin AFB

- National leaders in munitions development and predictions
- Large explosive test facilities
- Extensive collaborations with UF

We are not dependent on collaborations, but will seek to maximize them for broad impact

CM³C

Demonstration Problem: Free-Standing Mound Configuration

We have verified the free-standing mound configuration in static tests

Raw Image **Contour Recognition**

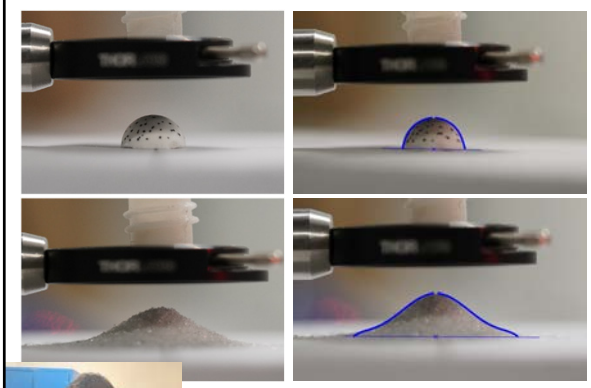

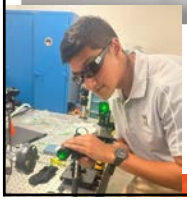

Work performed by Navi Singh
US Army Westpoint Cadet
and Summer 2025 research intern

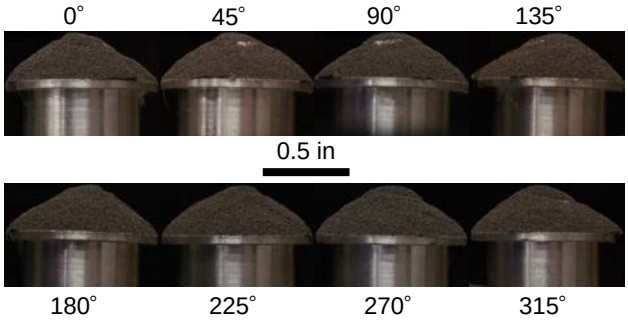

Takeaway: Possible but likely problematic for larger, outdoor shots

Concerns include (i) uneven coating, (ii) shot-to-shot variability, (iii) time on range needed to coat the charges, and (iv) disruption by breezes before initiation

Demonstration Problem: Unconfined Hemispherical Shell



Nunes *et al*. (2025) AIAA SciTech demonstrated use of Kinetic Sand™ as an unconfined hemispherical shell

We have now advanced to custom shells using pre-characterized particles

- 100 nm flake, 800 nm spherical, and 10 um spherical SiO_2 particles
- Polydimethylsiloxane (PDMS) binder (silicon-based organic polymer)
- Size distributions quantified in advance using a Malvern Mastersizer





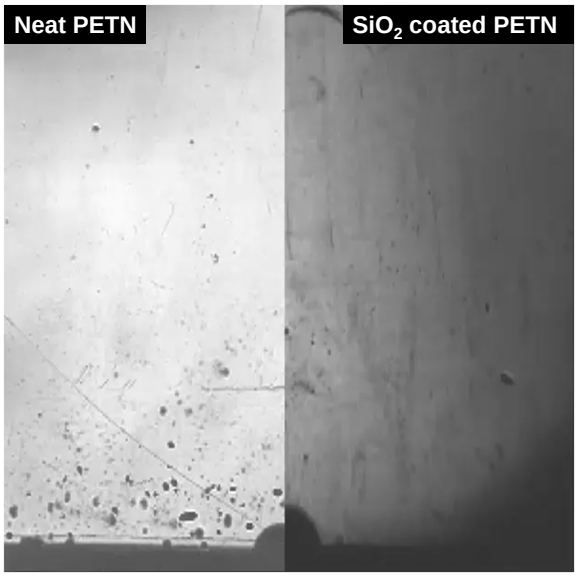

Schlieren imaging of the bare and coated explosive

CM³C This work was partially supported by a 2023-2025 SNL Academic Alliance LDRD led by A. Hammond-Clements

Demonstration Problem: Extinction Diagnostics

Diffuse Back-Illuminated Extinction Imaging (DBI-EI) Quantifies Light Extinction, $\tau = -\ln(I/I_0)$

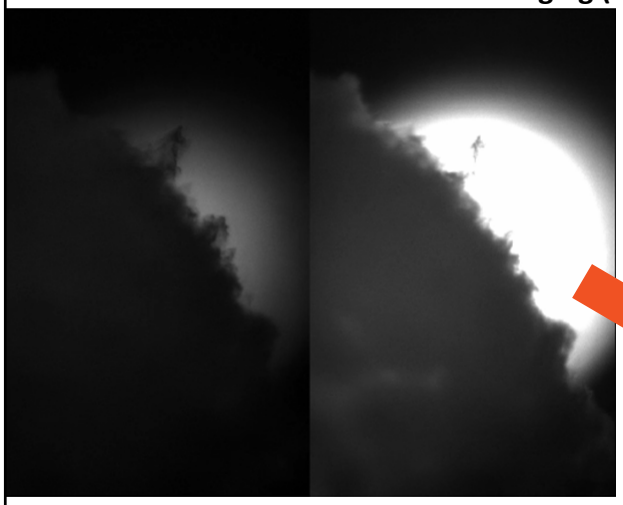

DBI-EI of the 100 nm SiO2 coated PETN hemisphere

Extinction is a function of particle concentration and governed by fundamental Mie scattering

- Accurate measurements must also account for particle size, index-of-refraction, optical configuration, and multi-scattering. See Nunes *et al*. (2026) *Optics Express.*

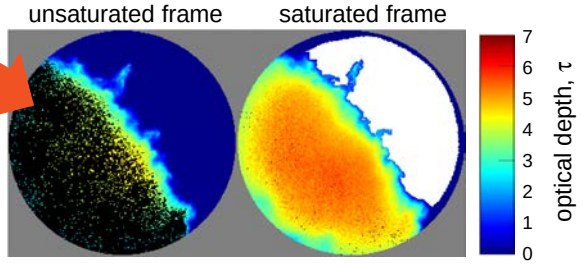

Extinction measurement from an example frame

CM³C This work was partially supported by a 2023-2025 SNL Academic Alliance LDRD led by A. Hammond-Clements

Demonstration Problem: Extinction Diagnostics

Line-of-sight extinction (bold line) is ensemble and azimuthally averaged from five repeat shots (dotted lines)

Uncertainty calculated assuming uncorrelated uncertainty in measured quantities

$$u_\tau^2 = \left(\frac{\partial \tau}{\partial I}\right)^2 u_I^2 + \left(\frac{\partial \tau}{\partial I_0}\right)^2 u_{I_0}^2$$

where $u_I = u_{I0} = I_N$: camera noise floor

$$u_\tau^2 = \left(\frac{I_N}{I_0}\right)^2 \left(e^{2\tau} + 1\right)$$

- As expected, uncertainty is the greatest where the extinction, τ, is the highest

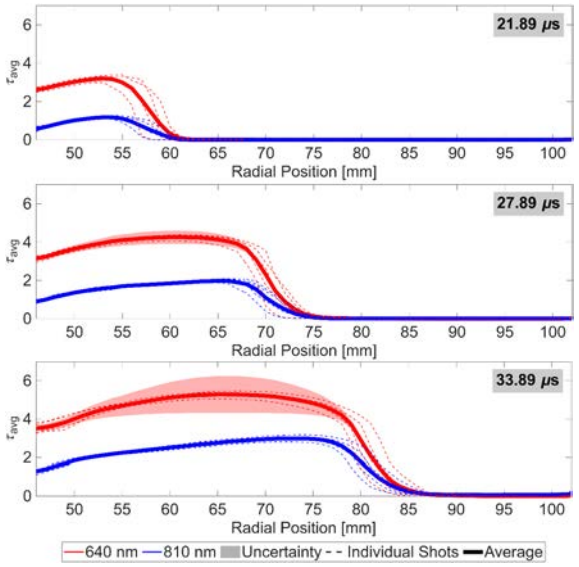

CM³C This work was partially supported by a 2023-2025 SNL Academic Alliance LDRD led by A. Hammond-Clements

Demonstration Problem: Extinction Diagnostics

Ensembled average line of sight extinction is

$$\bar{\tau}(r) = \int_{-\infty}^{\infty} N(s) C_{\text{ext}}(s)\, ds$$

were s is the pathlength, N is local number concentration, and C_{ext} is the extinction cross-section

Deconvolution yields the radially resolved extinction coefficient ($\beta_{\text{ext}} = N \cdot C_{\text{ext}}$)

- Results use the onion peeling method of Dash (1992) Appl. Opt.
- Uncertainty estimated from Monte Carlo variation of τ

Results were presented at the 2026 meeting of the Central States Section of the Combustion Institute, and a journal manuscript is in progress.

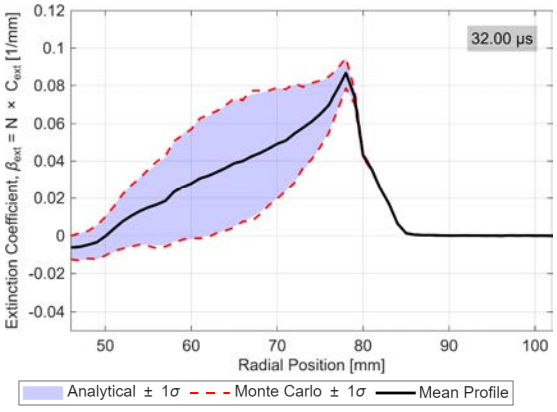

We have developed a systematic measurement framework for particle concentration

CM³C This work was partially supported by a 2023-2025 SNL Academic Alliance LDRD led by A. Hammond-Clements

Demonstration Problem: Reactive Particle Shots

First Reactive Particle Test Campaign Performed April 2026 at Purdue

12 mm diameter PETN charges, coated with Mg/PDMS and C/PDMS shells tested in outdoor shots

- **QoI1:** high-speed schlieren imaging for blast wave locations
- **QoI2:** pencil gages for pressure
- **QoI3:** visible extinction for particle concentration

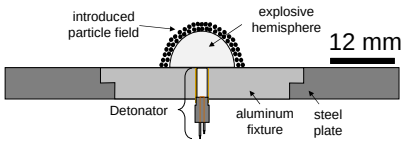

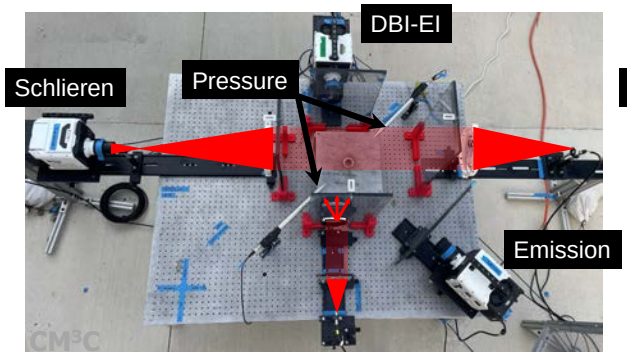

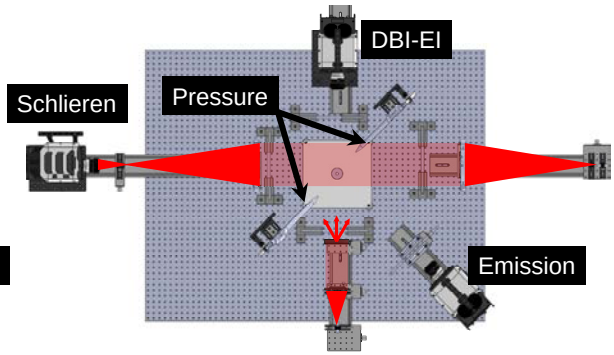

CM³C

Demonstration Problem: Reactive Particle Shots

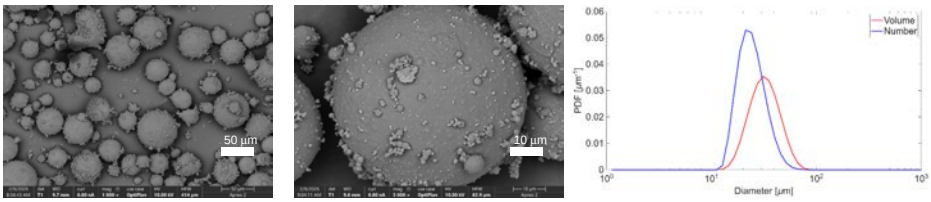

Nominally 45 µm Magnesium Particles Purchased from US Nanomaterials

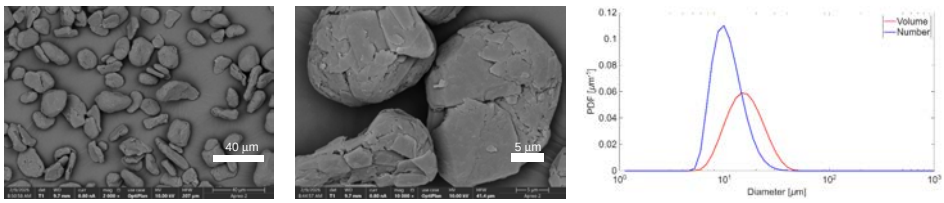

Nominally 25 µm “Spherical” Carbon Particles Purchased from US Nanomaterials

CM³C

Demonstration Problem: Reactive Particle Shots

Coatings proved to be inconsistent in the field

- Predictions of reactions due to the added binder may also be challenging

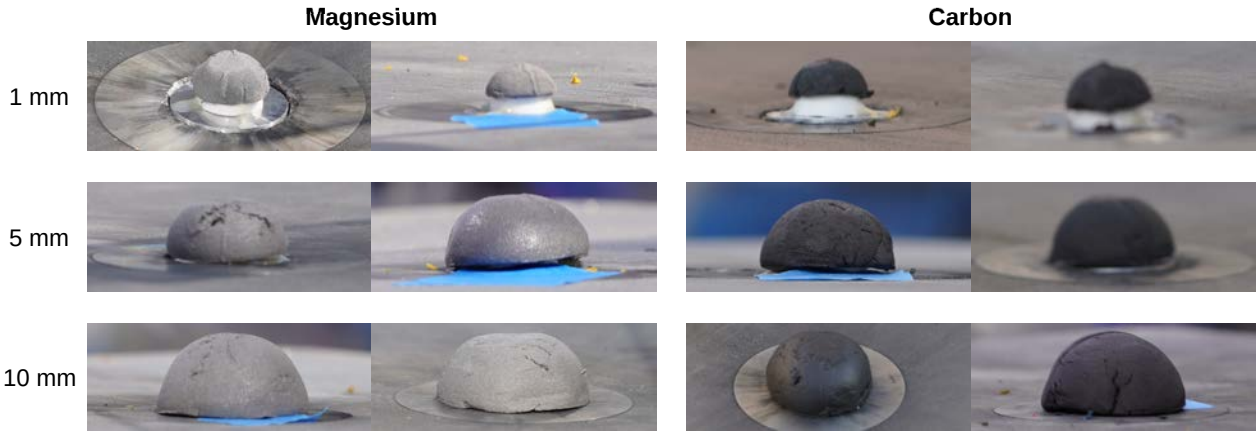

CM³C

Demonstration Problem: Reactive Particle Shots

As expected, thicker shells impede expansion rates at early times

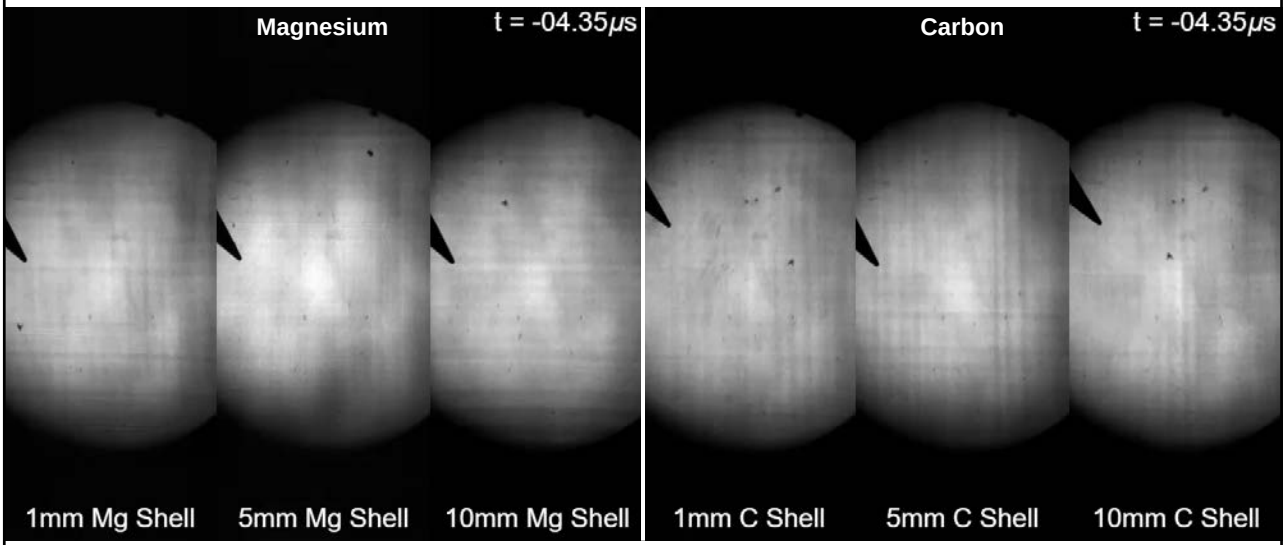

Demonstration Problem: Reactive Particle Shots

As is often the case, visible emissions are complex to interpret.

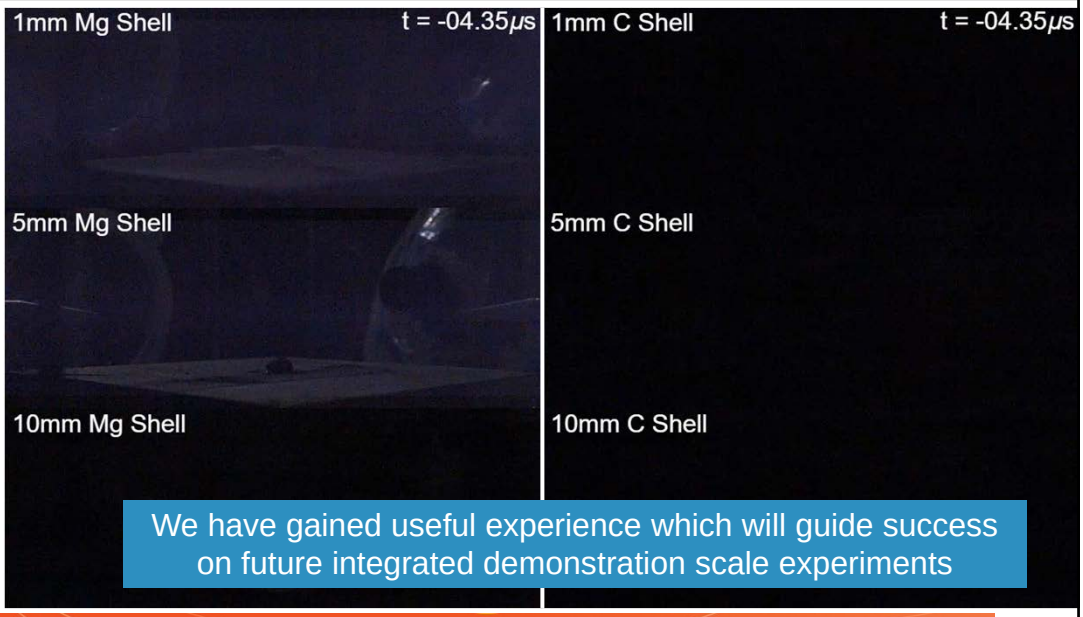

We have gained useful experience which will guide success on future integrated demonstration scale experiments

CM³C

Demonstration Problem: Reactive Particle Shots

Intent is to do Larger-Scale Experiments at a Partner Organization, hopefully in FY27

We are currently actively discussing with:

- ***NSWC Crane Indiana***: POC Eric Scheid. They have both indoor and outdoor test arenas and experienced explosive operators. They are currently reviewing the feasibility of PETN pressing and financial ability to support their team.
- ***AFRL Eglin Florida***: POC Don Littrell. They supported the UF PSAAP-II center and likely the nations best blast characterization arena. We have submitted a white paper for AFRL support of PSAAP-IV shots.
- ***SNL Albuquerque***: POC Ryan Marinis. SNL has been conducting similar hemisphere experiments stemming from prior collaborations with this team. Performed initial tests in Summer 2026 with a PSAAP-IV student intern (Cohen Nunes), but it is not yet decided if they can support particle shots.
- ***LLNL***: POC Robert Reeves. Performing hemispherical shots similar to SNL. Again, a collaborative test campaign may be possible.

NAVSEA WARFARE CENTERS CRANE

Sandia National Laboratories

Lawrence Livermore National Laboratory

CM³C

Demonstration Problem: Reactive Particle Shots

At the larger-scale, we hope to reduce uncertainty in the particle shell

- No PDMS binder with poorly known chemistry
- Larger charge reduces relative influence of initiation spot size
- Thicker shell such that most particles are not crushed by the shock waves
- Desired diagnostics: (i) high-speed imaging, (ii) pressure gauges, (iii) x-ray or visible extinction
- We further hope to engage Chris Murzyn of SNL with his blast-hardened laser absorption gauges, likely leveraging current JMP funding

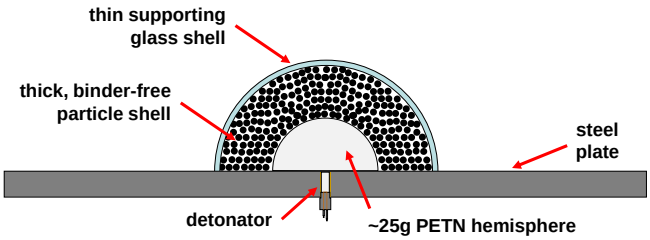

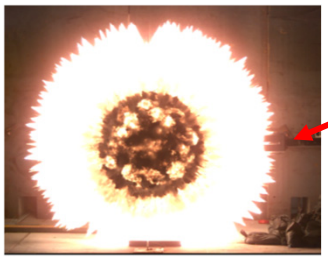

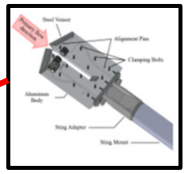

Murzyn *et al.* (2024) *J. Appl. Phys*

CM³C

CM³C Microscale Problem

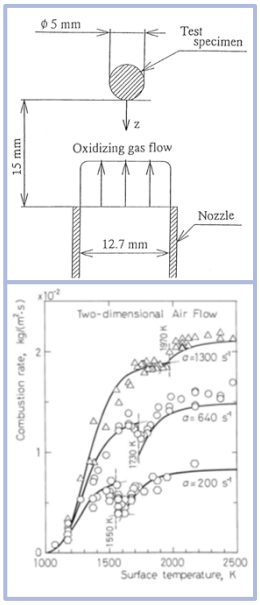

The combustion of carbon rods has been extensively studied by Makino *et al.* 1990-2011

- 5 mm Ø rods, Joule heated to ignition
- Surface temperatures measured by radiography
- Regression rates measured by videography

- Results have been used to validate HF-PR simulations of isolated particles →

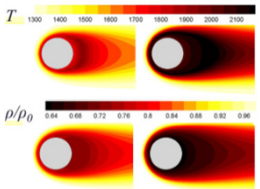

Jiang *et al.* (2022) *Int. Commu. Heat Mass*

Again, proven experiments with a foundation of existing data

We plan to acquire new experimental data for several V&V objectives:

1. Validate HF-PR simulations of *group combustion of particle arrays*
2. Quantify uncertainty in chemical kinetics and validate HF-PR modeling assumptions (such as negligible radiation effects)
3. Support uncertainty reduction via computational informed experiments (relay validation)

CM³C Microscale Problem

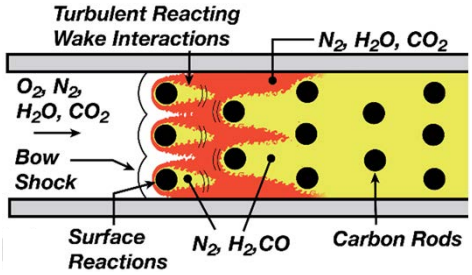

Conditions: Flow velocities up to Mach 1.5, 300 K static temperature (500 K stagnation temperature)

- Equivalent to leading particles in a multiphase blast propagating into the surrounding air

Diagnostics: (i) High-speed videography of regression rates, (ii) IR pyrometry of surface temperature, (iii) Laser absorption spectroscopy (LAS) of gas products, (iv) Laser induced florescence (LIF) of intermediates

Enabled by the world leading high enthalpy infrastructure at the Maurice J. Zucrow Laboratories



8 lbs/sec, 1500 F heated air system

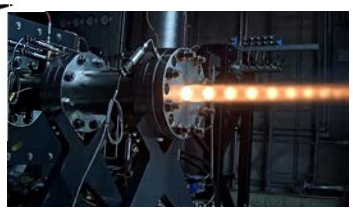

ramjet testing at Zucrow

CM³C

CM³C Microscale Problem

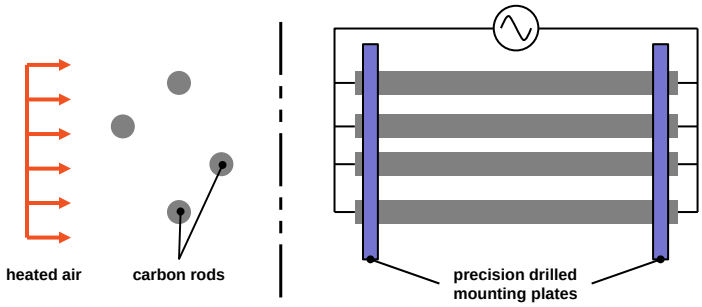

Experimental Concept:

- Carbon rods mounted in precision drilled plates for flexible 2D configurations
- Electrically heated to ignition
- Side-view imaging for straight forward measures of diameter changes

Experimental Plans:

1. First experiments will validate single rod measurements against published data
2. Quantify clusters of several closely-spaced rods that induce wake-wake and group combustion dynamics
3. Investigate geometries and quantities from our large collection of HF-PR simulations which most significantly impact the uncertainty budget (relay validation)

CM³C

Microscale Problem: Single Rod Results

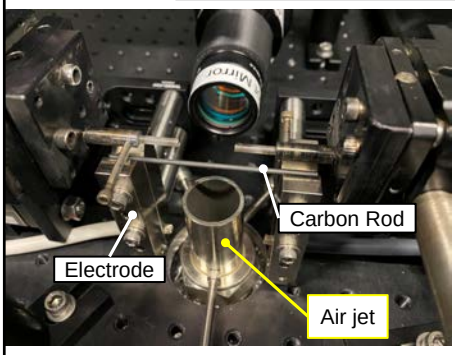

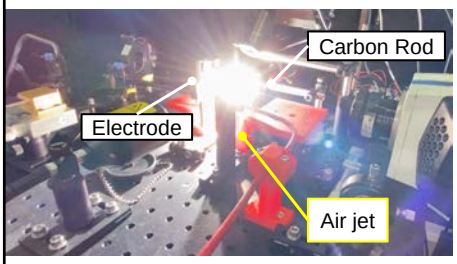

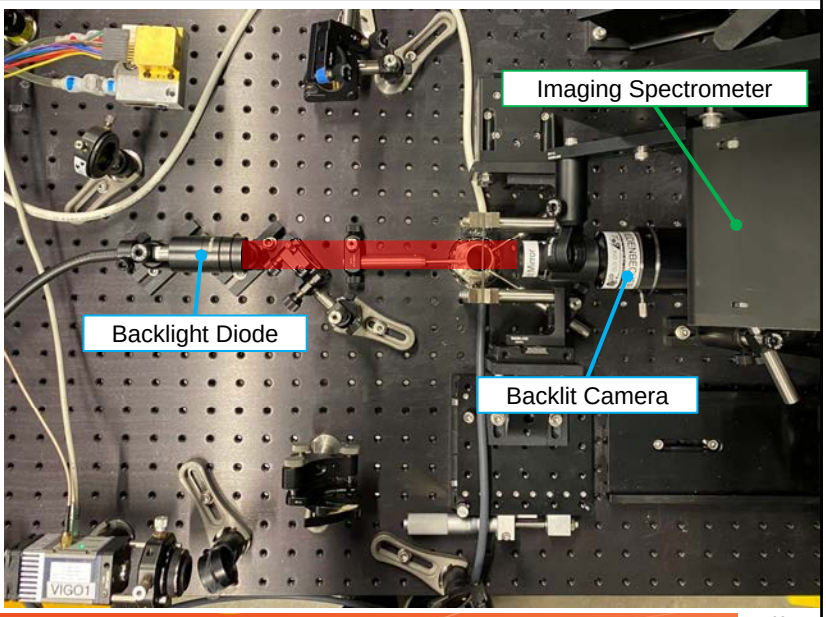

Microscale Problem: Single Rod Results



500 µm

Airflow Direction

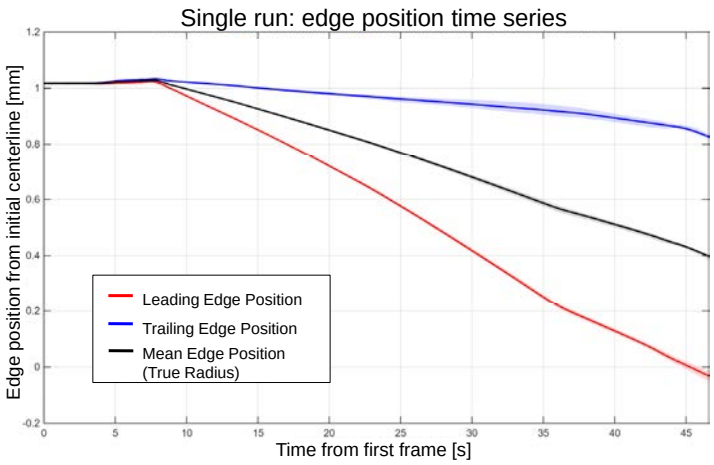

Combustion measured by tracking leading edge and trailing edge positions over time relative to the initial centerline. Bands show 90% variation across field of view.

Microscale Problem: Single Rod Results

Accurate temperature measurements have been a bit more challenging than expected

- We have tried two-color pyrometry, IR imaging, and imaging spectroscopy
- Imaging spectroscopy provides both spatial and spectral resolution; therefore, the highest confidence

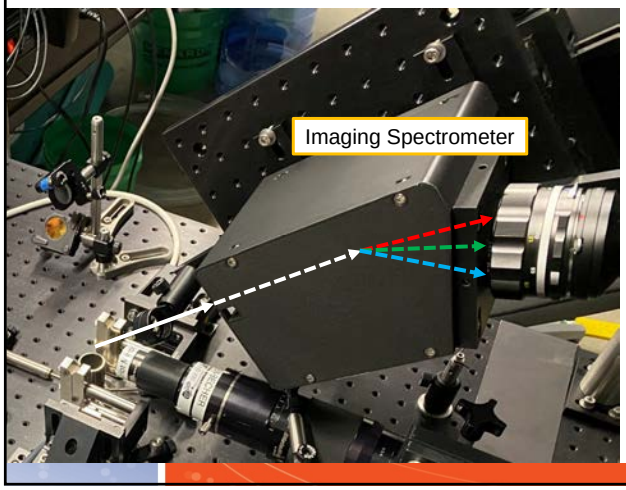

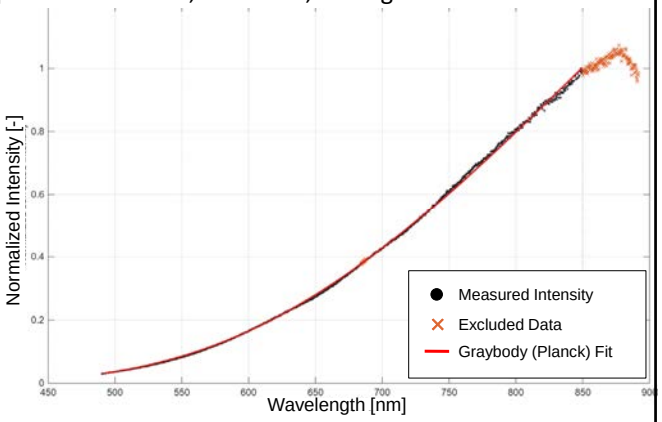

Sample graybody fit for a single row of a single frame
$T = 1990K, R^2 = 0.9996, MSE = 3.704 \times 10^{-5}$

Microscale Problem: Single Rod Results

Consistent with prior results, combustion onsets at $T \gtrsim 1000$ K

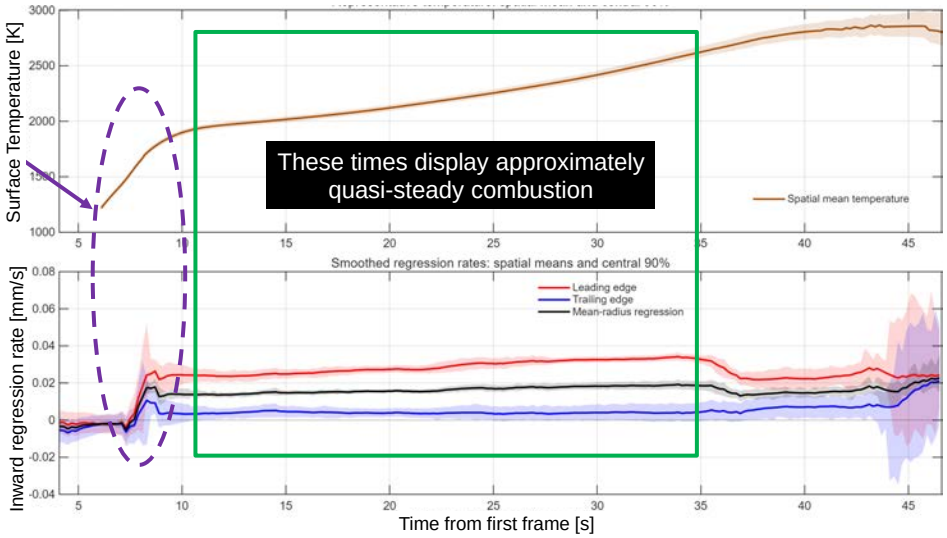

Temperature along with the leading, trailing, and mean regression rates over time. Bands show 90% variation across field of view.

CM³C

Microscale Problem: Single Rod Results

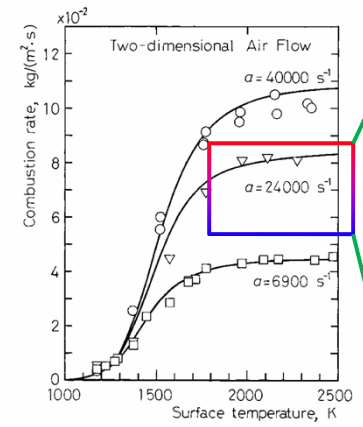

Prior experiments (symbols) and model (—) by Makino *et al*.

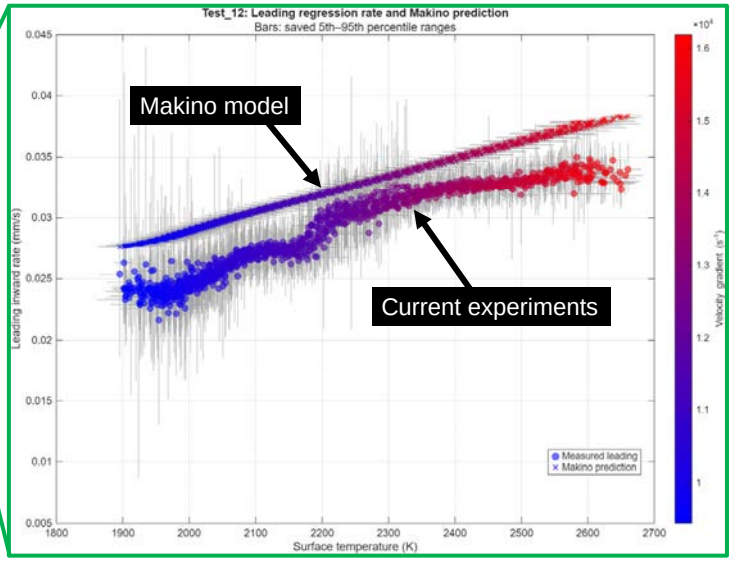

Current experimental results compared to the Makino *et al*. model

CM³C

Microscale Problem: Single Rod Results

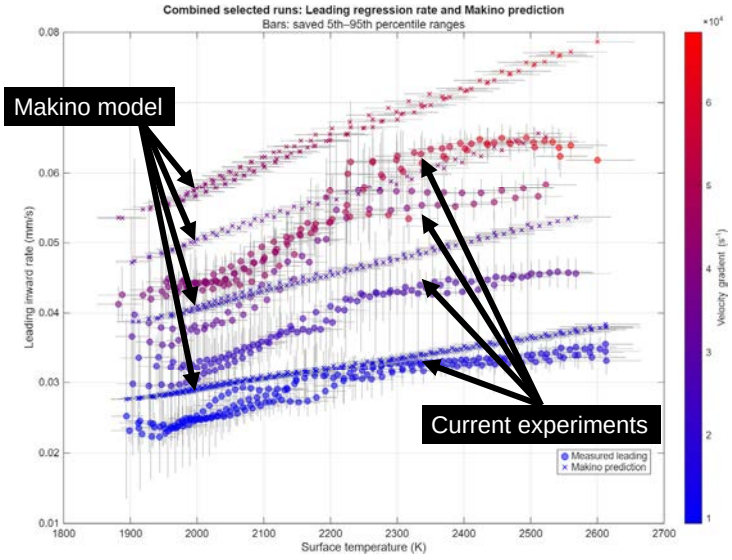

Current experimental results compared to the Makino *et al*. model

Results show more systematic disagreement with prior model at increasing air flow rates. This is indicative of either:

1. Unresolved experimental biases (E.g., we are currently investigating local air velocities.)
2. Deficiencies in prior models or the experiments to which they are calibrated

Carbon rod experiments have been advanced to high-quality measurement results

Microscale Problem: Next Steps

Laser Absorption Spectroscopy is quantifying CO concentration and temperature in the near field

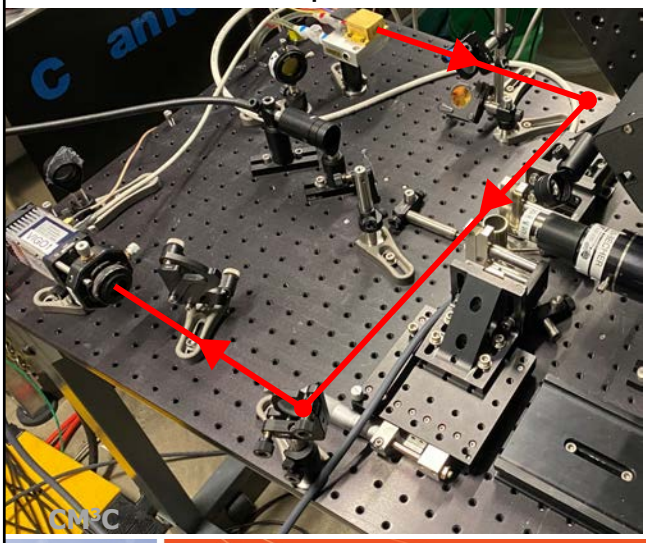

We are progressing to a multi-rod configuration with the next step being two rods.

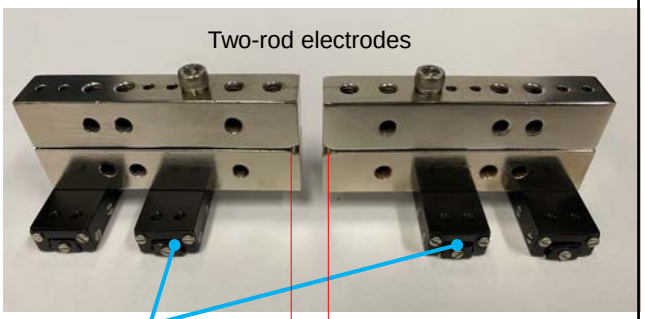

Rails accommodate axial thermal expansion separately for each rod.

Arbitrary Rod Spacing

Single rod electrode



CM³C Mesoscale Problem

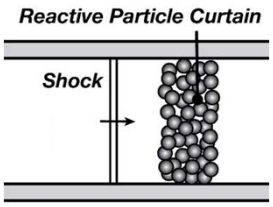

Mesoscale experiments will validate simulation of :

- Dense particle combustion
- Turbulence and transport
- Clustering/voids

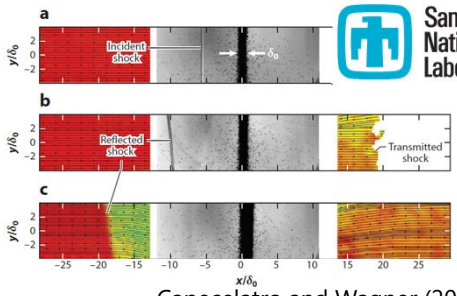

Capecelatro and Wagner (2024)
Annu. Rev. Fluid Mech.

Proposed experiments will build on current efforts to include:

- Reactive curtains of carbon and magnesium particles
- Post-shock temperatures >2500 K in the SNL High-Temperature Shocktube (HST) will ignite particles in a strong convective flow
- Imaging, LAS, CARS, PIV, and other diagnostics that leverage a long history of PU and SNL partnerships
- A novel curved curtain would further induce fingering instabilities

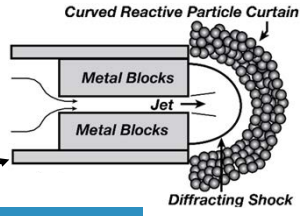

Collaborations will further our long track record of joint experiments and student exchange

CM³C

CM³C Mesoscale Problem

This team also has a long-history of experimental collaboration with LANL, LLNL, and other NNSA sites

- Example: P-Rad measures of explosive particle dispersal
  - Proposed experiments will consider reactive C and Mg particles
- Broad partnership opportunities for (i) reactive ejecta, (ii) carbon ablatives, (iii) explosive device testing, (iv) and other NNSA missions

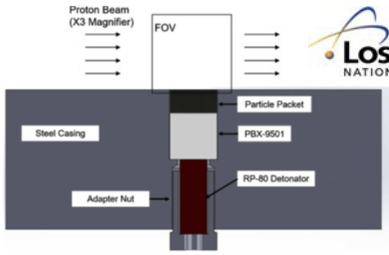

Hughes, Balachandar, *et al.* (2021), *Shock Waves*

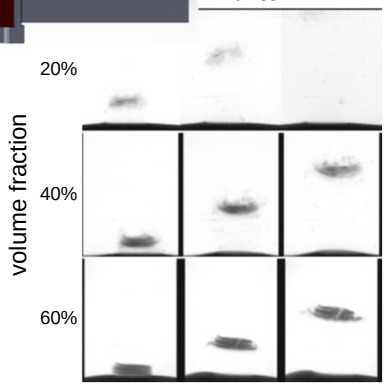

We are not dependent on collaborations, but will seek to maximize them for broad impact

Mesoscale Problem: Progress To Date

We have made the least progress here.

- Hoped for LDRD ideas and JMP proposals at SNL have not yet gained traction
  - We continue to work with SNL partners (e.g., Kyle Daniel, Justin Wagner, Daniel Richardson) for upcoming proposal opportunities
- Purdue PhD candidate, Cohen Nunes, completed his 10-week PSAAP internship at SNL in summer 2026 where he contributed to related aluminum shock tube experiments (and explosive tests).

PURDUE UNIVERSITY

Summary

Experiments have largely achieved the year 1 goals and are progressing into planned year 2 activates.

- Current pain point/risks:
  1. We continue to debate the best way to construct the demonstration charges
  2. Mesoscale problem in not currently supported at out lab partners
  3. Efforts will transition to new students following the upcoming graduation of a few senior students

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

Verification, Validation, and Uncertainty Quantification, Part 2

Co-PI: Nam-Ho Kim

Mechanical and Aerospace Engineering

University of Florida

NNSA National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

CM³C V&V UQ and Workflow Team



Nam-Ho Kim
(Florida)



James Fairbanks
(Florida)



Jacob Stewart
(Florida)



Isaly Tappan
(Florida)



Raul Valle
(Florida)



Alex Millauer-Hsu
(Purdue)



Cohen Nunes
(Purdue)



Gregory O'Brian
(Purdue)

QoIs and Key Sources of Uncertainty

Basic QoIs

- QoI1: Blast wave locations as a function of time
- QoI2: Time-dependent pressure and impulse at selected radial distances
- QoI3: Volume fraction of reactive particulate at selected radial locations

Higher-order metrics

- QoI4: Temperature & heat release rate
- QoI5: Numbers & amplitude of fingering instabilities
- QoI6: Time-dependent PDFs of pressure and particulate volume fraction



Dominant Sources of Uncertainty

- Stochastic uncertainties (aleatory) and errors
  - Initial random distribution of particle sizes, locations, and density
  - Stochastic detonation characteristics (PETN)
  - Stochastic turbulent mixing
- Experimental measurement uncertainties and errors
  - Shock structures and strength, particle concentrations, pressure history, temperature and species concentration
  - Measurement processing uncertainty and error
- Modeling & numerical uncertainties and errors
  - Spatially filtered gas-phase governing equations
  - Reduced-order surface & gas-phase chemical reactions (C and Mg)
  - ML surrogate model uncertainty
  - Radiation and compaction model uncertainty (minimally contributing)

CM³C

Uncertainty Budget for Uncertainty Reduction (UR)

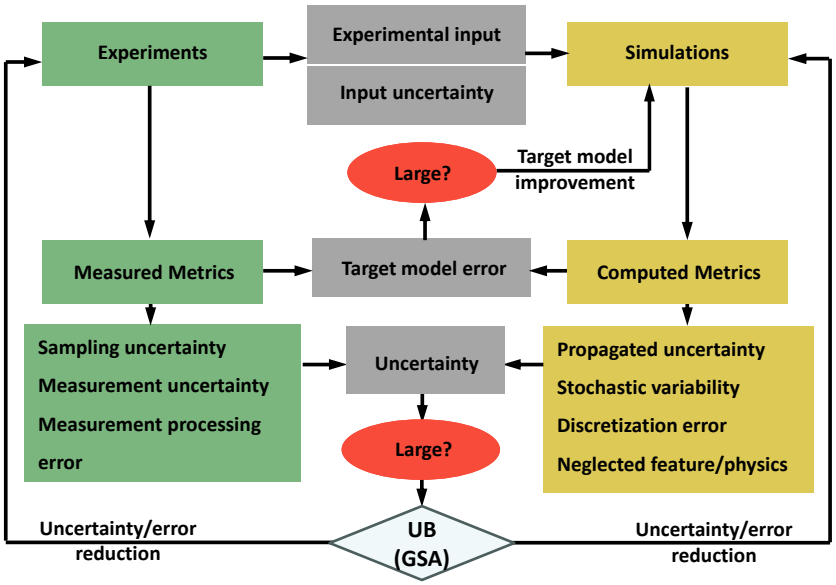

CM³C

Uncertainty Quantification, Propagation & Reduction

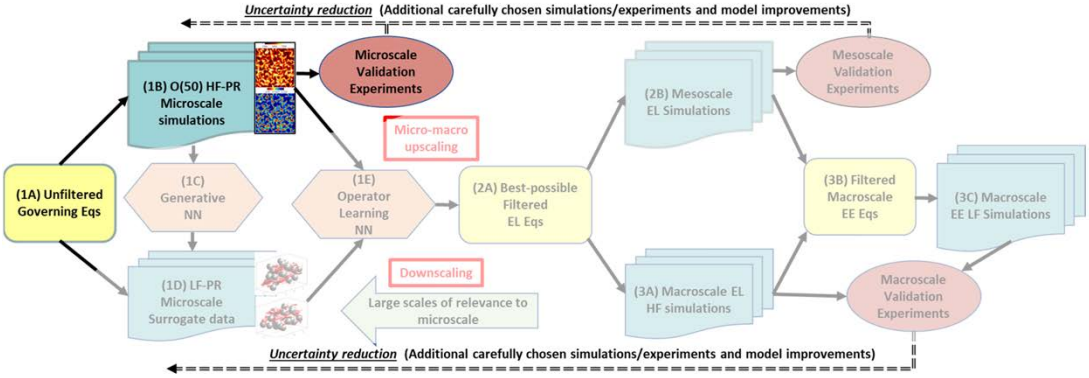

- Exascale HF-PR Simulations (1B) → Almost full resolution
- Epistemic uncertainty from chemistry & radiation
- Unrepresented physics (particle shape, porosity, and other imperfections)
  - ➢ Purpose of the microscale rod campaign
- Insufficient resolution at higher Re
  - ➢ Grid resolution study

CM³C

Uncertainty Quantification, Propagation & Reduction

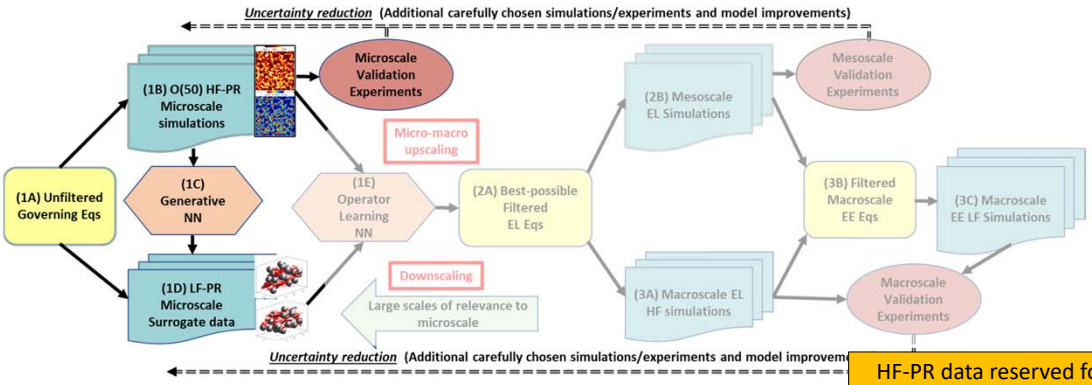

HF-PR data reserved for testing will serve as benchmark for UQ

- LF-PR Surrogates (1D) generated using the Generative NN (1c)
- Propagated uncertainty is the same as in HF-PR simulations (in the limit of a perfect surrogate)
  - Insufficient training (limited HF-PR data)
    - ➢ Vary data size and k-fold cross validation and study convergence
  - Epistemic uncertainty from Generative NN architecture limitations & training protocol
    - ➢ Consider different architectural parameters and establish uncertainty

CM³C

Uncertainty Quantification, Propagation & Reduction

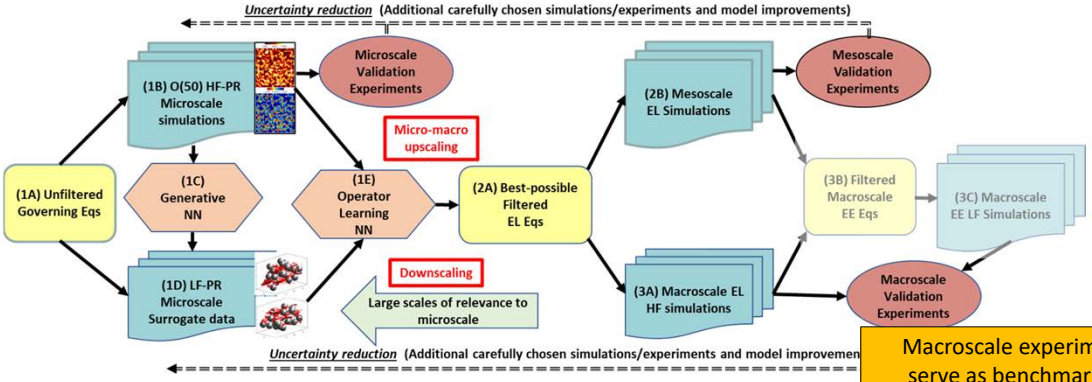

Macroscale experiments will serve as benchmark for UQ

- HF-PR + LF-PR data → OLNN → Best-possible EL Eqns → Macroscale EL simulations
  - Propagated uncertainty is bounded by HF-PR and LF-PR (in the limit of a perfect closure)
  - Insufficient training (limited HF-PR and LF-PR data)
    - ➢ Vary data size and k-fold cross validation and study convergence
  - Epistemic uncertainty from OLNN architecture limitations & training protocol
    - ➢ Consider different architectural parameters and establish uncertainty

CM³C

Hierarchical UQ Framework

- Relay validation at different stages
- Uncertainty propagation/reduction from a lower scale to an upper scale
  - Microscale UQ of reduced surface and gas-phase chemical reactions
  - Mesoscale UQ of OLNN closure models and group combustion effect
- Propagating both error and uncertainty
- Uncertainty reduction to expose model error (Bayesian method)

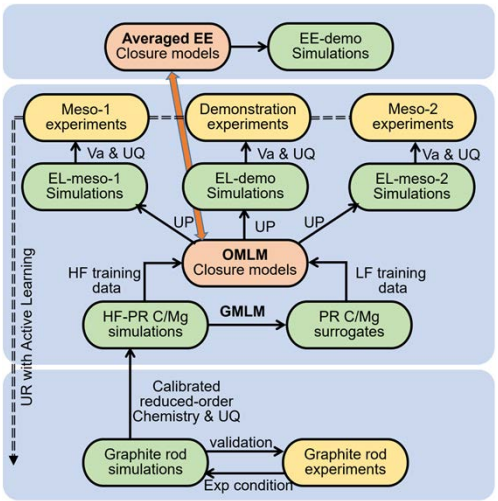

CM³C

UF UNIVERSITY of FLORIDA

Microscale UQ

- Experiment: Combustion of carbon rod(s) within a high-speed (Ma 1.5) airflow (temp. 300°K)
  - Measurements: Regression rate, surface temperature, gas composition
- Simulation: Eulerian diffuse interface simulation
  - unfiltered gas and particulate phase governing equations in fixed domain
  - heterogeneous surface and gas-phase reactions (optimal reduced-order reactions)
- Objective: UQ of reduced surface and gas-phase chemical reactions
- Initially with a single rod, later multi-rods (guided by simulations)

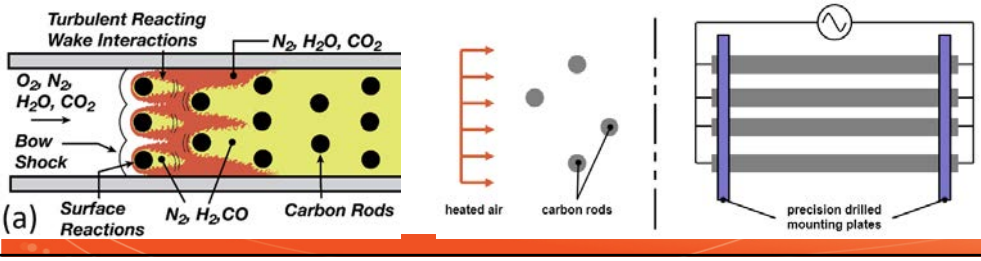

CM³C

UF UNIVERSITY of FLORIDA

Microscale UQ

- Uncertainty in experiments
  - Inputs: Electric current, humidity, air flow profile, rod properties
  - Outputs: regression rate, surface temperature (spectroscopy), chemical compositions (LAS)
- Uncertainty in simulations
  - Inputs: propagated uncertainty, discretization uncertainty, Carbon surface and gas chemistry (burning model)
  - Outputs: regression rate, chemical compositions
- Discrepancies between experiments and simulations
  - Experiments control electric current, while simulations specify surface temperature (experimental measurements become simulation inputs)
  - Fixed rod radius (simulation) vs. variable rod radius (experiment)

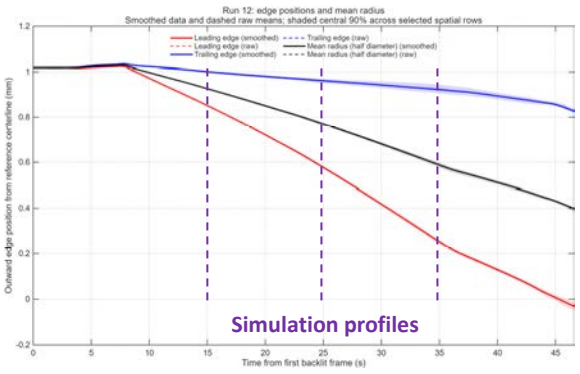

CM³C

Measurement Processing Uncertainty

- Major source of epistemic uncertainty in experiments
  - Commonly happened when QoI cannot be measured directly
  - Rigorous UQ and UR are required to have a meaningful validation
- One-rod Eulerian diffuse interface simulation to UQ in measurements
  - Not only measurement variability but also measurement processing uncertainty

Intensity calibration of Spectrometer

Low exposure lamp

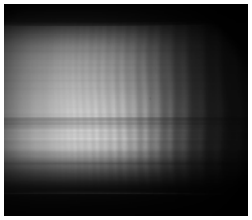

High exposure lamp

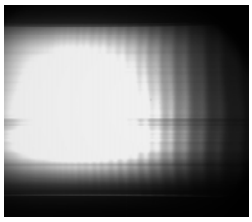

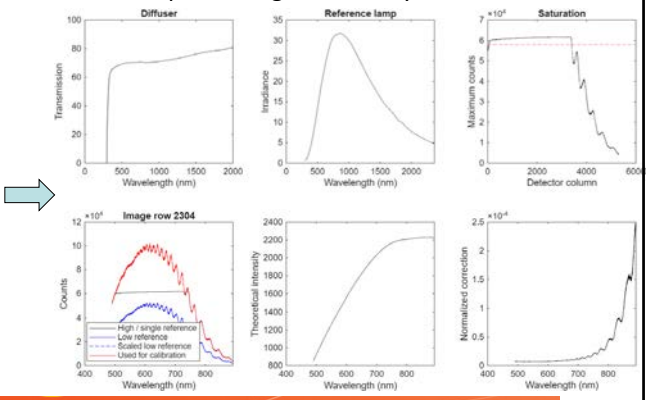

CM³C

Relay Validation

- Design/reduce epistemic uncertainty in experiments using simulation
  - Simulation will decide position of two rods to maximize the difference in regression rates and minimize the radiation effect
  - Where to place the measurement probes for maximum sensitivity

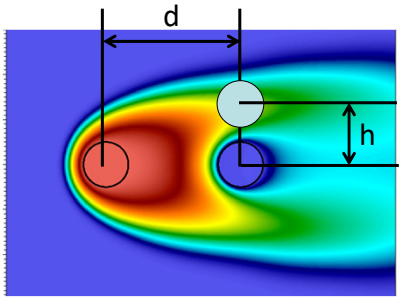



Two-rod electrodes with arbitrary spacing

CM³C

Sources of Uncertainty in Machine Learning Models

- Data uncertainty (aleatory uncertainty)
  - Test data itself contains the uncertainty (UQ of data will be available)
- Model-form uncertainty (Epistemic uncertainty)
  - Network architecture
- Training uncertainty
  - Uncertainty in construction of training and validation dataset.
  - Initialization of parameter (local minima)

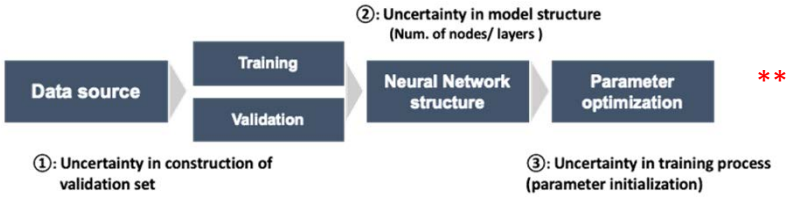

Divide and conquer

CM³C

ML Training Uncertainty

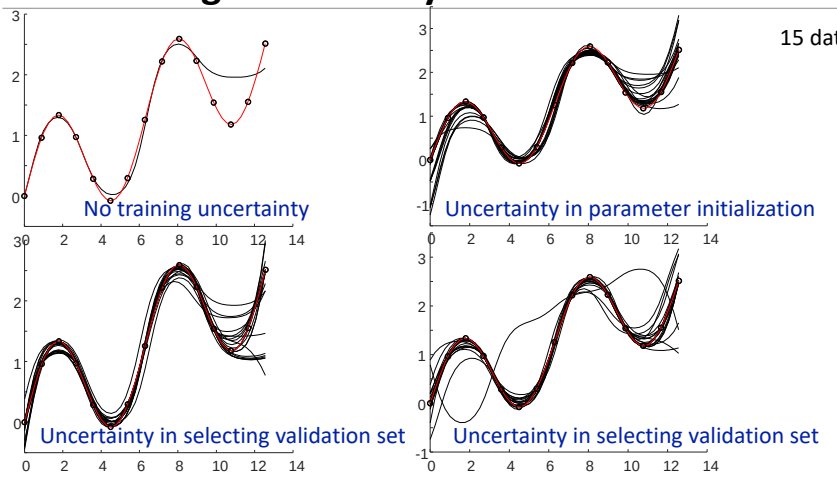

15 data from $z = 0.2x + \sin(x)$

Source of uncertainty	Mean of e_{RMSE}	Standard deviation of e_{RMSE}
No uncertainty	0.2798	0.0
Initialization uncertainty	0.1900	0.1216
Training set uncertainty	0.1888	0.1106
Both uncertainties	0.1930	0.2258

Only achievable with repeated trainings

CM³C

UQ Using Linear Perturbation

- Goal: How to estimate uncertainty using a single dataset and single training?
- We view ML as a surrogate model $f(\mathbf{x}, \mathbf{b})$ with nonlinear regression
- Nonlinear regression also minimizes error (loss function)

$$e_i = y_i - f(\mathbf{x}_i, \mathbf{b}), \quad i = 1, \ldots, n_y$$

$$\mathbf{b}^* = \underset{\mathbf{b}}{\operatorname{argmin}}\ e_{RMS}$$

- Once parameters are trained, we linearize the model (w.r.t. $\mathbf{b}$) at the trained point

$$f(\mathbf{x}, \mathbf{b}) \simeq f(\mathbf{x}, \mathbf{b}^*) + \sum_i \frac{\partial f(\mathbf{x}, \mathbf{b}^*)}{\partial b_i} \Delta b_i$$

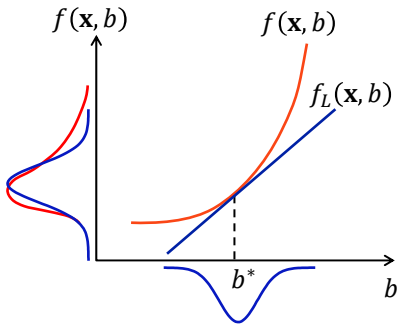

Linear function for $\Delta\mathbf{b}$ and $\partial f / \partial b_i$ as coefficients

UQ Using Linear Perturbation *cont.*

- Theories from linear regression provide useful information for UQ
- Uncertainty in samples & model form → uncertainty in parameters & prediction

1. Sampling uncertainty: $\sigma^2 = \frac{SSE}{n}$
2. Variance of parameters: $\mathbf{\Sigma_b} = \sigma^2[\mathbf{X}^T\mathbf{X}]^{-1}$
3. Prediction variance: $\sigma_y^2 = \hat{\sigma}^2 \boldsymbol{\xi}(\boldsymbol{x}_p)^T (\mathbf{X}^T\mathbf{X})^{-1} \boldsymbol{\xi}(\boldsymbol{x}_p)$

$$\mathbf{X} = \begin{bmatrix} \frac{\partial f(\mathbf{x}_1, \mathbf{b}^*)}{\partial b_1} & \frac{\partial f(\mathbf{x}_1, \mathbf{b}^*)}{\partial b_2} & \cdots & \frac{\partial f(\mathbf{x}_1, \mathbf{b}^*)}{\partial b_{n_\beta}} \\ \frac{\partial f(\mathbf{x}_2, \mathbf{b}^*)}{\partial b_1} & \frac{\partial f(\mathbf{x}_2, \mathbf{b}^*)}{\partial b_2} & \cdots & \frac{\partial f(\mathbf{x}_2, \mathbf{b}^*)}{\partial b_{n_\beta}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f\left(\mathbf{x}_{n_y}, \mathbf{b}^*\right)}{\partial b_1} & \frac{\partial f\left(\mathbf{x}_{n_y}, \mathbf{b}^*\right)}{\partial b_2} & \cdots & \frac{\partial f\left(\mathbf{x}_{n_y}, \mathbf{b}^*\right)}{\partial b_{n_\beta}} \end{bmatrix}$$

Design matrix:
Available as a part of the training process

UQ Using Linear Perturbation (Test Function)

- Samples from a simple function and assess the accuracy and uncertainty of ML model
- Surrogate Function: $Ae^{-kt} + c$
- Neural Network Fit: $f(t) = \sum_h V_h \tanh(w_h t + b_h) + c)$
- Compared STD between repetitions and linear perturbation

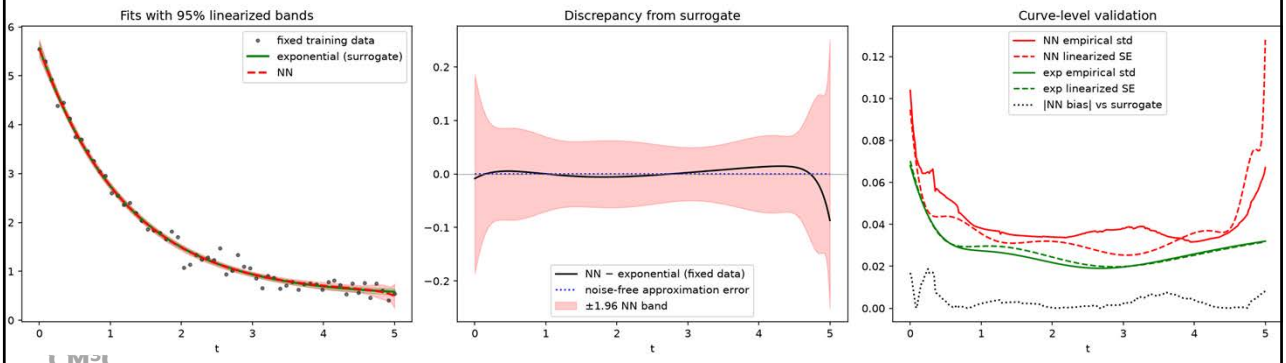

UQ of Training Uncertainty

- Uncertainty due to randomly selected training/validation sets

1. We view early stopping criterion as K-fold cross-validation process
2. In linear regression, K-fold cross-validation error can be estimated using one training

$$e_{PRESS} = \sqrt{\frac{1}{n_y - 1}\sum_{i=1}^{n_y}\left(\frac{e_i}{1 - E_{ii}}\right)^2}$$

$$\hat{\mathbf{y}} = \mathbf{X} \cdot \mathbf{b} = \mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{y} = \mathbf{E} \cdot \mathbf{y}$$



UQ of Model-Form (Network Architecture) Uncertainty

- Stepwise backward elimination
  - In ML, hidden nodes play the role of bases in function approximation
- UQ using linear perturbation can estimate uncertainty in weights

$$\sigma_{b_i} = \hat{\sigma}\sqrt{[(\mathbf{X}^T\mathbf{X})^{-1}]_{ii}} \qquad c_i = \frac{\sigma_{b_i}}{|b_i|}$$

Estimated STD of weight b_i — CoV

- Remove the hidden node(s) with the most uncertainty weight(s) (insignificant)
  - Repeat the process until the R_a^2 increases and/or e_{PRESS} increases

This process requires repeated trainings!

Uncertainty Reduction in ML Models

- Prediction variance of linearized mode: $\sigma_y^2 = \hat{\sigma}^2\boldsymbol{\xi}(\boldsymbol{x}_p)^T(\mathbf{X}^T\mathbf{X})^{-1}\boldsymbol{\xi}(\boldsymbol{x}_p)$
- GSA to separate epistemic uncertainty from aleatory uncertainty
- Active learning strategy to add samples when $\sigma_{\text{epistemic}} \geq \sigma_{\text{aleatory}}$
  - Add samples at maximum $\sigma_y^2(\boldsymbol{x}_p)$

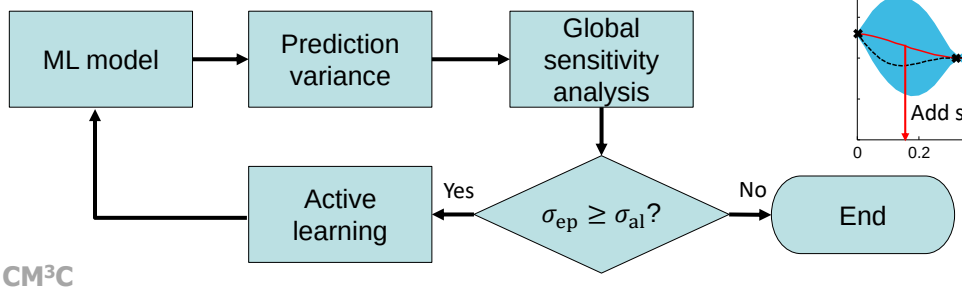

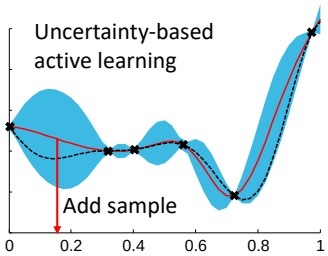

CM³C

Questions?



NNSA
National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

MFC for Particle Resolved Simulations

Daniel J. Vickers

School of Computational Science and Engineering, GT

Advised by Spencer Bryngelson

Georgia Tech | UF UNIVERSITY of FLORIDA | PURDUE UNIVERSITY

NNSA National Nuclear Security Administration

Georgia Tech

Particle-Laden Flows in Detonations

- Particle debris is of interest in simulations of detonations
  - Particles shape the direction of the flow
  - The flow entrains particles
- Particle—fluid interactions can produce complex structures
  - "Particle fingers" form clusters in high-speed channels
  - Particle distribution gaps create high-speed gas jets

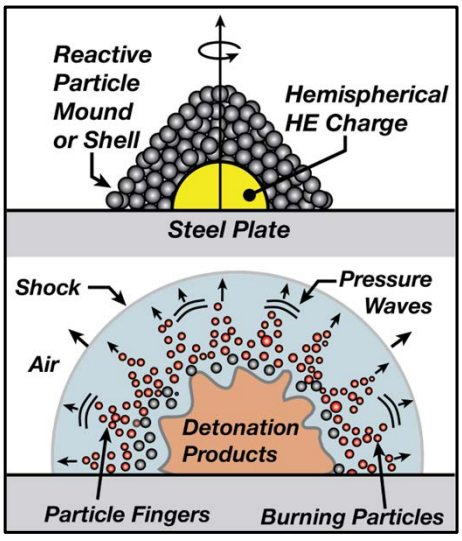

Where Do Immersed Boundaries Fit In?

- Immersed boundary optimizations improve the particle-resolved microscale simulations

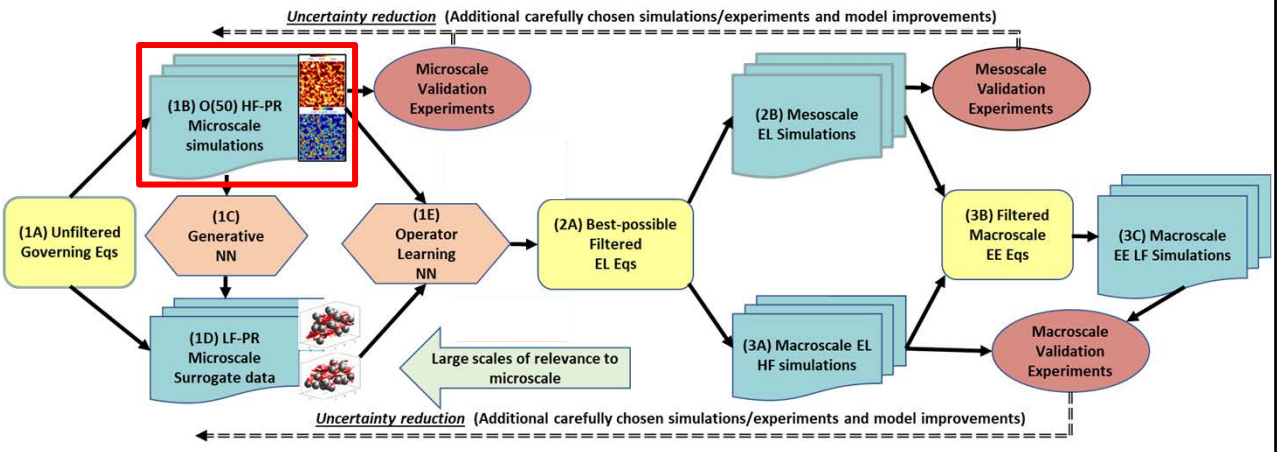

CM³C

Software Status in April

- Last TST meeting, we demonstrate capability for a 1,000-particle shock curtain
- Features shown
  - 2-way coupling for shock—particle interactions
  - Collisions
  - Supported GPU and MPI integration
- Run on
  - 800x200x200 grid
  - 4 NVIDIA A100 GPU for 2 hours
- Full-system scaling was not demonstrated

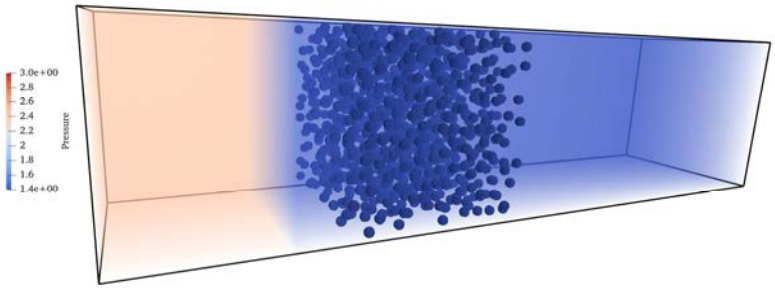

CM³C

State of the Code

- Current IBM features include
  1. Rajat Mittal type 2nd order IBM
  2. Soft-sphere collision model
  3. Periodic boundary support
  4. Reacting particle surfaces
  5. Particle distributions can be generated parametrically
- Performance Improvements
  1. Optimized ghost-cell identification
  2. Optimized collision detection

Costs of PR Simulations

- Since April, I have focused on scaling and stability
- To utilize exascale machines, efficient scaling is required
  1. Shown is a scaling efficiency plot of the current largest PR simulation published
     a. 8,192 Particles per rank
     b. 8.4 million particles total
     c. This is on the order of 10's of TB
  2. As parallel efficiency decreases, the computational cost of these simulations grows

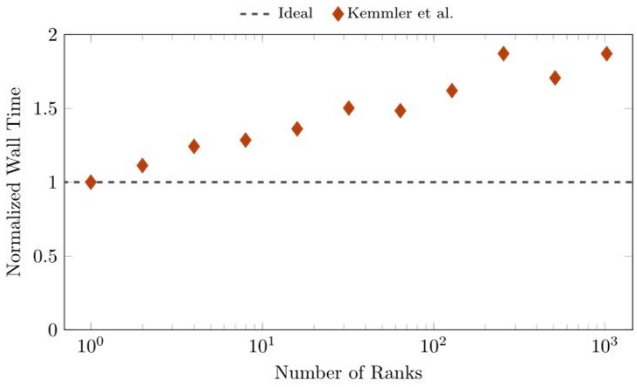

"Efficiency and scalability of fully-resolved fluid-particle simulations on heterogeneous CPU-GPU Architectures" – Kemmler et al. IJHPCA (2023)

Weak Scaling in Tests in MFC

- We tested the MFC IBM weak scaling
  - Scaled to the full system of OLCF Frontier
  - 16,384 particles per GCD
    - 10 grid cells per diameter
  - 360^3 grid cells per GCD
    - Particle count limited by number of grid cells
  - Simulation inviscid for example purposes
- Our base case is 64 GCDs
  - Uses inter-node communication in all directions
  - Increased size 8x until full machine was reached

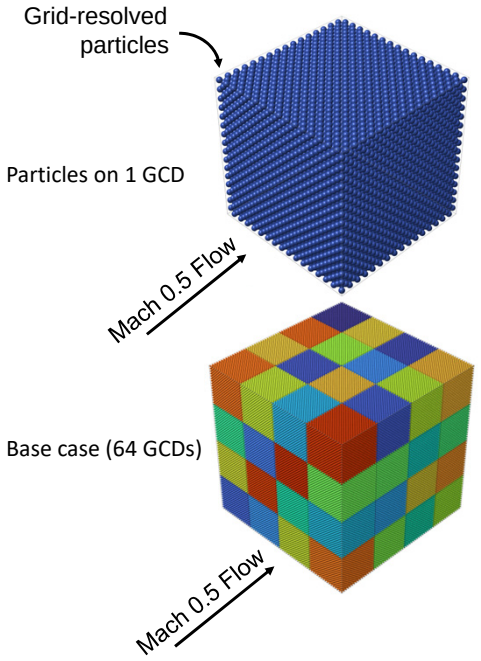

Scaling Achieved in MFC

- This case was scaled to 96% of Frontier
  1. At 32,768 ranks we measured 1.036 wall time (96.5% efficient)
  2. At 74,088 ranks we measured 1.128 wall time (88.7% efficient)
     1. This simulation contained 1.2 billion grid-resolve particles

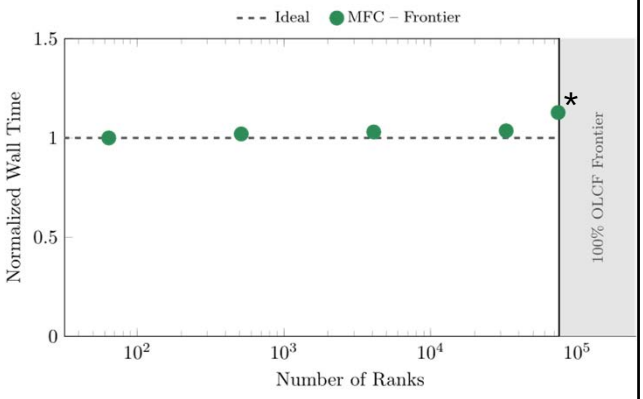

* Caveat: Different simulation campaign, unclear result

Scaling Achieved in MFC

- This case was scaled to 96% of Frontier
  1. 32,768 GCDs: 97% efficiency
  2. 74,088 GCDs: 89% efficiency
     1. This simulation contained 1.2 billion grid-resolve particles
- The Kemmler data renormalized to 64 ranks
  1. We scaled with 2x particles per rank
  2. The Kemmler reference is an LBM solver
  3. We scaled to 144x the number of particles
- At 512 ranks
  1. MFC: 98% efficiency
  2. Kemmler: 90% efficiency

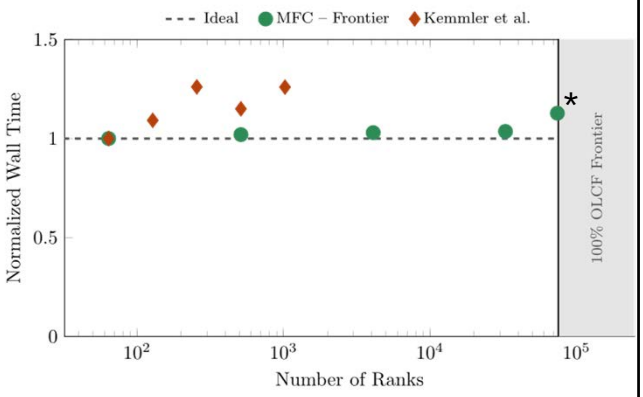

* Caveat: Different simulation campaign, unclear result

Movie

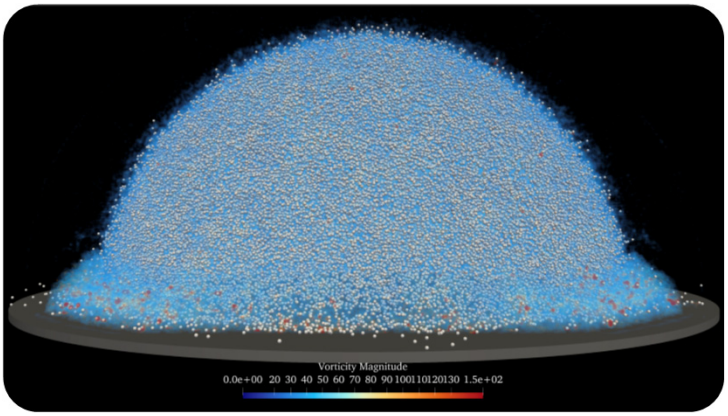

Conclusion

- New features
  - 2nd order IBM
  - Reacting surfaces
  - Stability improvements
- Scaling
  - 1.2 billion particles: 89% efficiency
- Hemispherical Detonation
  - 100,000 magnesium particles
  - 20 grid-cell diameters
  - Mach 25 shock

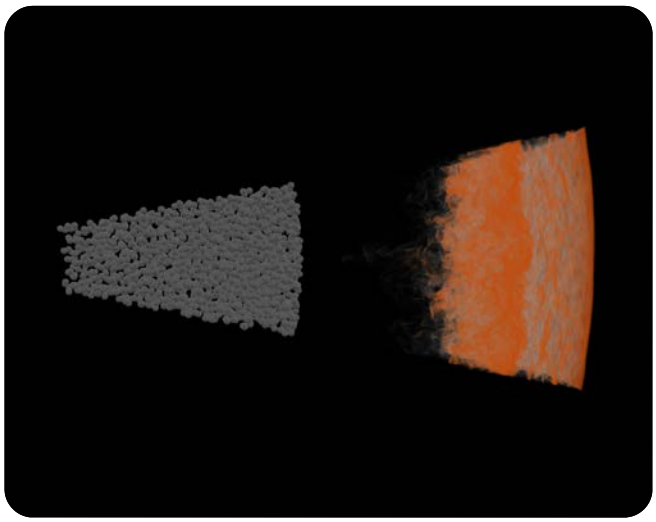

CM³C

CM³C

Questions?

Backup

Scaling Progress in MFC

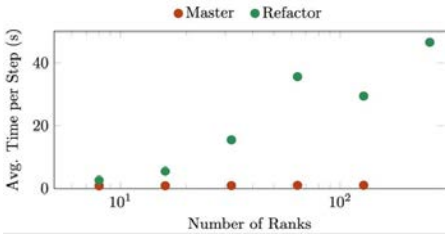

5/4/26

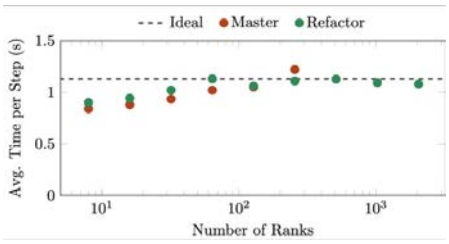

5/11/26

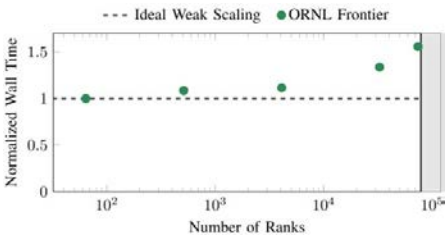

7/13/26

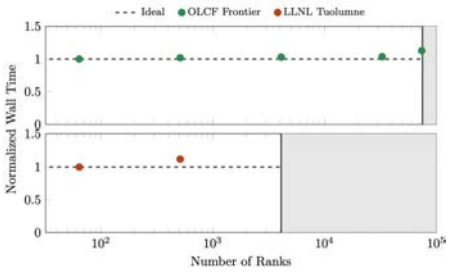

10/5/26

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

MFC for Euler-Lagrange Simulations

T. Daoud, B. Wilfong, J. Aguilar, J. Gillis

Department of Mechanical and Aerospace Engineering

NNSA National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

UF UNIVERSITY of FLORIDA

Euler-Lagrange Simulations

- Conduct the best-possible EL meso- and macro-scale simulations
- Include as much physics in record-breaking EL simulations
- Comparative runs against micro- and macro-scale
- Exascale simulations on thousands of GPU ranks

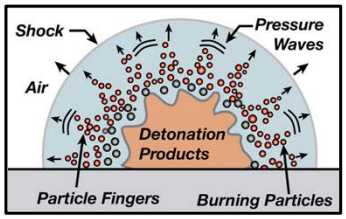

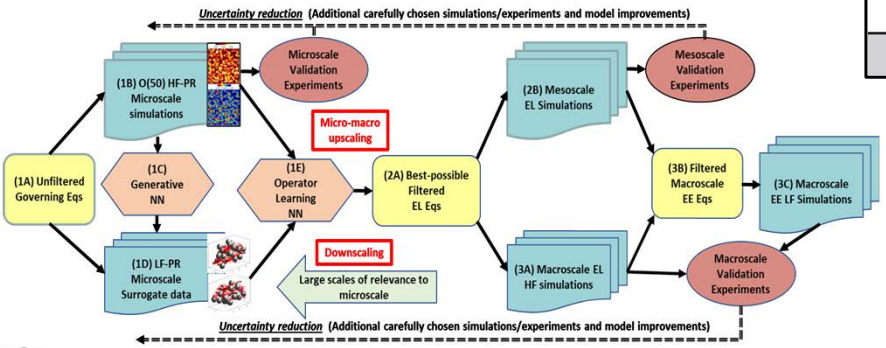

CM³C

Euler-Lagrange Framework in MFC

- MFC supported Lagrangian bubbles, but not solid particles
- Efforts have been invested into further developing this capability
- Validations are underway with the further inclusion of physics

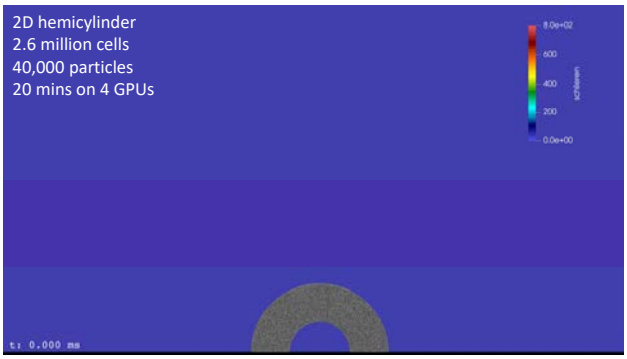

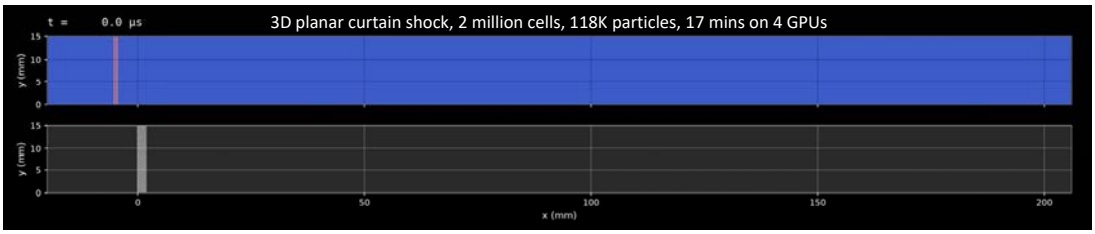

CM³C

Euler-Lagrange Framework in MFC

- We are working on a GPU-ready Euler–Lagrange particle solver, in several phases.
- Generalizable modules shared amongst Lagrangian solid particles and bubbles.
- Similar philosophy will be adopted for the contact forces between IBM and EL.

Phase 1	✓ Force models (QS, QS Fluct, AM, PG) ✓ Two-way coupling (interpolation and projection) ✓ 26 regression tests, and various examples ✓ Memory allocation optimization
Phase 2	o Further performance optimization o Contact forces (collisional forces, torques)
Phase 3	o Particle load balancing
Phase 4	TBD

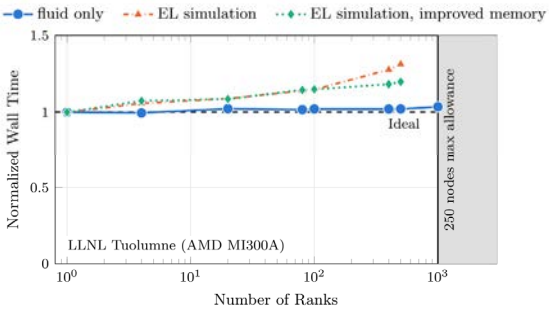

CM³C

Hemispherical Shell – Problem Setup

- 3D hemispherical shell sitting on high-pressure air
- 40 AMD MI300A ranks were used on Tuolumne, for 15 hours

Particle shell of magnesium property

Case	Packing Fraction	Particle Diameter	Shell Thickness	Number of Particles
1	15%	50-micron	15 mm	78 million
2			0.55 mm	910,000

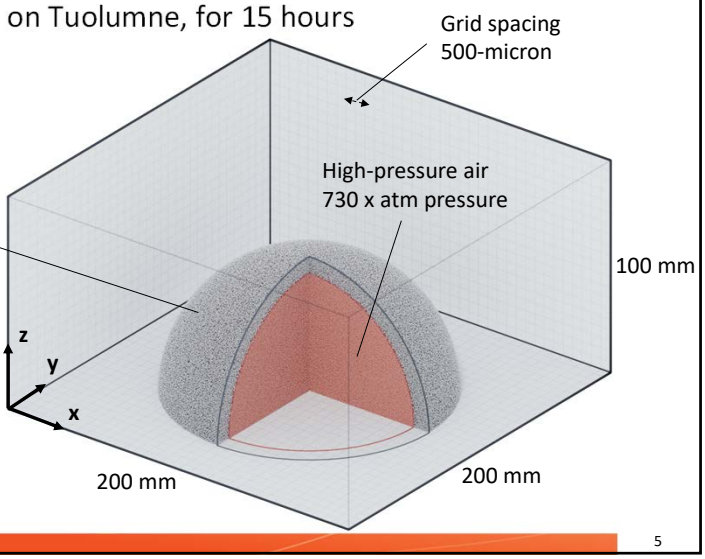

CM³C

Hemispherical Shell – **Case 1** – Animation

- 78 million particles surrounding a high-pressure region 730 x atm
- 40 AMD MI300A ranks were used on Tuolumne, for 15 hours

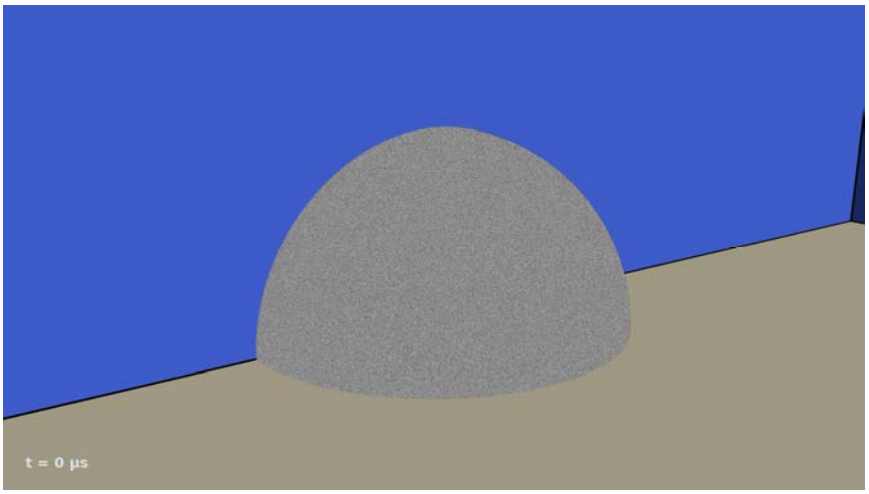

CM³C

Hemispherical Shell – **Case 1** – Plots

- 78 million particles surrounding a high-pressure region 730 x atm
- 2D slices portraying the schlieren at two different times

Where are the reflected shocks?



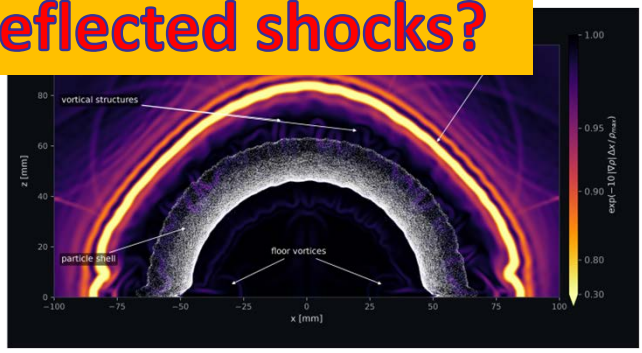

CM³C

Hemispherical Shell – **Case 2** – Animation

- 910,000 particles surrounding a high-pressure region 730 x atm
- 20 AMD MI300A ranks were used on Tuolumne, for 40 minutes

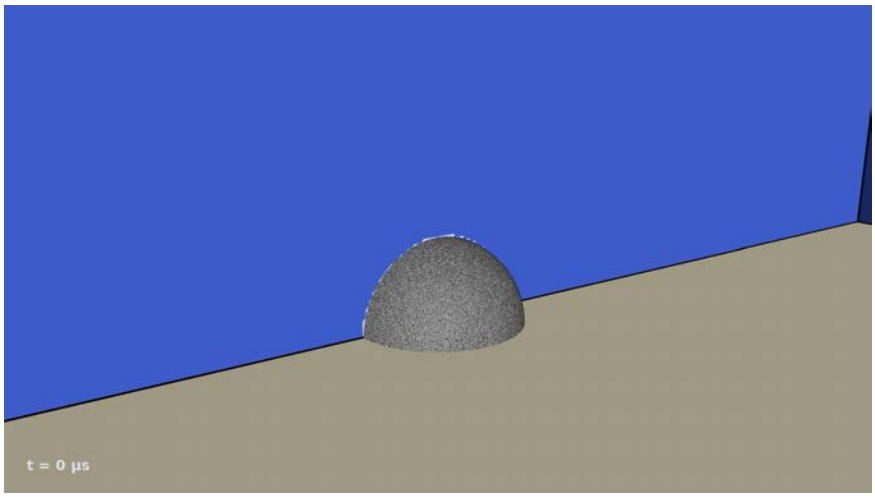

CM³C

Hemispherical Shell – **Comparison plots (1/2)**

- Schlieren comparison between cases 1 and 2 showing very different fluid behavior
- Reflected shocks are clearly observed in case 2, but none in case 1

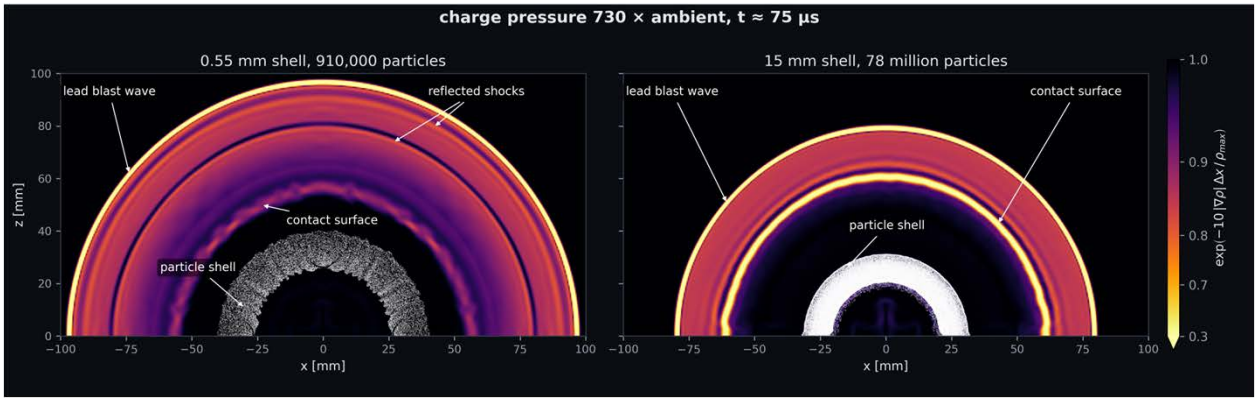

Hemispherical Shell – **Comparison plots (2/2)**

- Case 2 shows more instabilities across the contact surface due to repeated reflected-shock crossings
- Small-scale vortical structures, several times longer in case 1 than

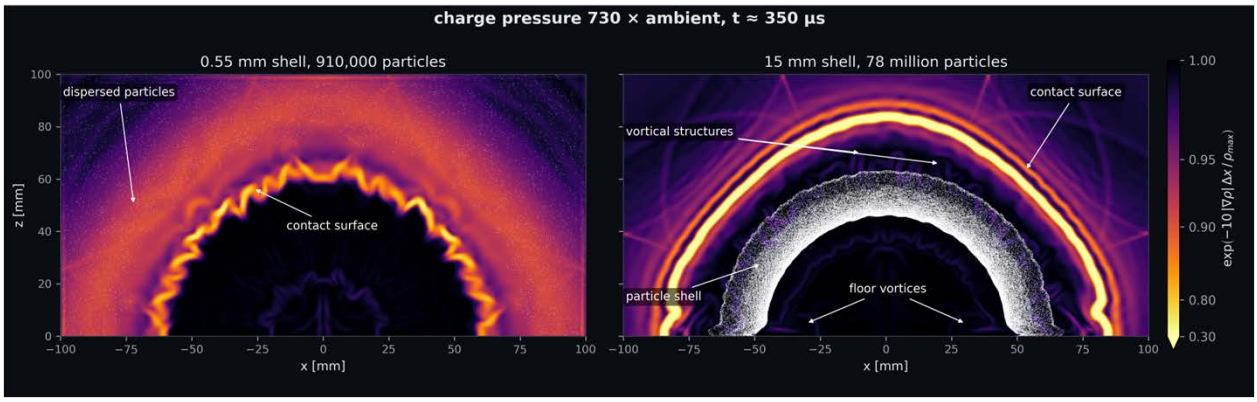

Hemispherical Shell – **Comparison plots (2/2)**

- Pressure plotted at the charge center shows the center implosions obtained in the thinner shell (case 2), whereas a piston-like effect in thicker shell (case 1)

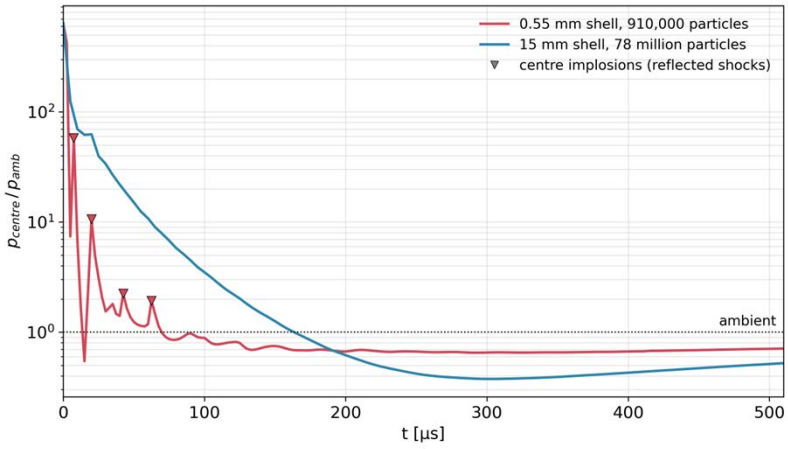

CM³C

Summer Internship at Lawrence Livermore National Laboratory

- Synthetic photonic doppler velocimetry numerical framework for particle laden flows

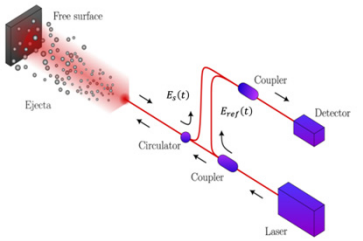

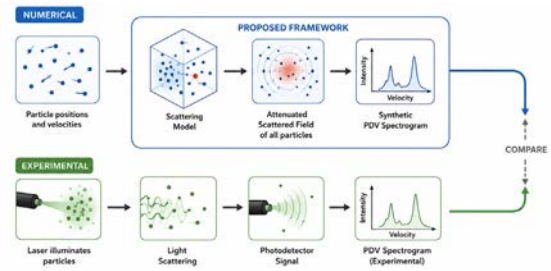



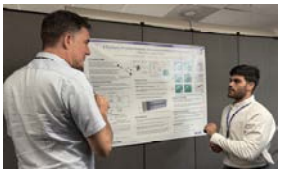

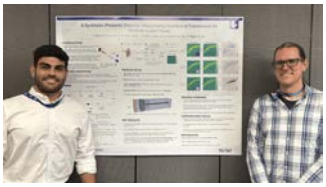

CM³C

Summer Internship at Lawrence Livermore National Laboratory

- Validations are underway against planar shock particle curtain data from LANL

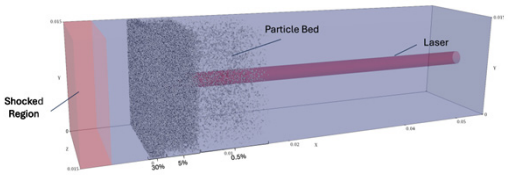

Synthetic PDV Spectrogram

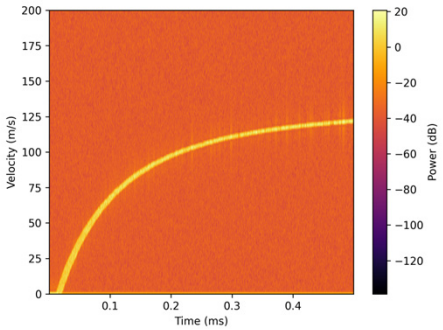

Spectrogram from LANL*

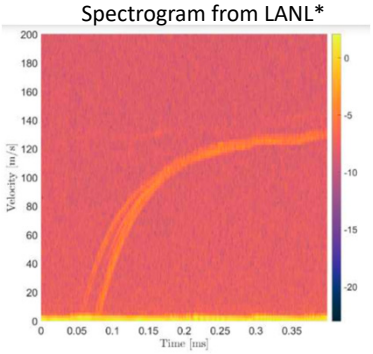

*Noble *et al.*, PDV Workshop, 2026.

CM³C

Summer Internship at Lawrence Livermore National Laboratory

- Potentially collaborate with Purdue over PDV diagnostic comparison on the hemispherical shell problem

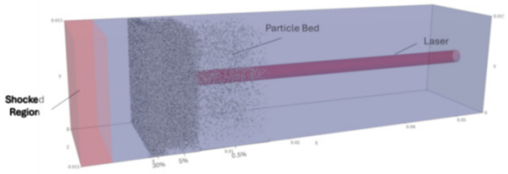

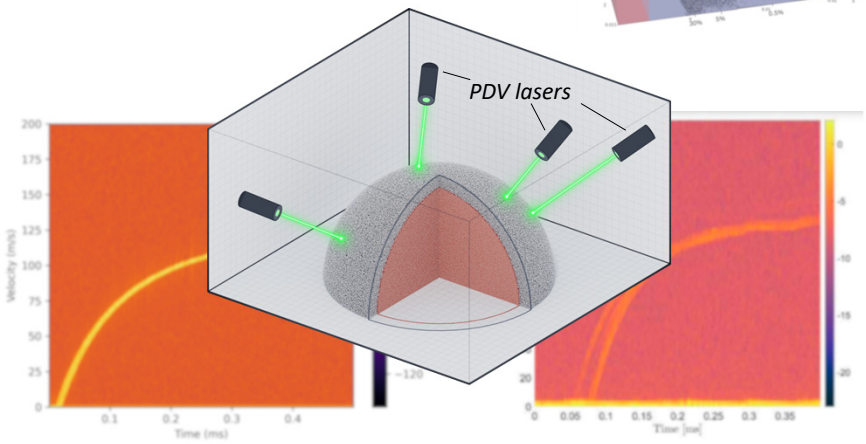

CM³C

CM³C





UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

CM³C: Center for Multiscale Modeling of Multiphase Combustion

T6: N-Particle Flow-Prediction: Theoretical Foundations and Rigorous Potential Flow Testing

Davin Anderson

Department of Mechanical and Aerospace Engineering

University of Florida

UF UNIVERSITY of FLORIDA · PURDUE UNIVERSITY · The University of Texas at Austin Oden Institute for Computational Engineering and Sciences

Introduction

- Current approaches to group combustion use averaged models to capture the neighboring effects
  - Decoupled from force/Reynolds stress models
- Our goal is to develop "best-possible" closure models that accurately account for neighbor interactions
- We are building a rigorous theoretical framework

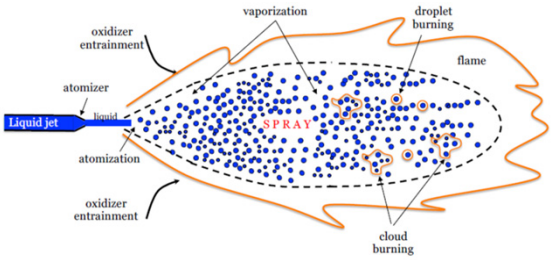

Overall Physics & Simulation Roadmap

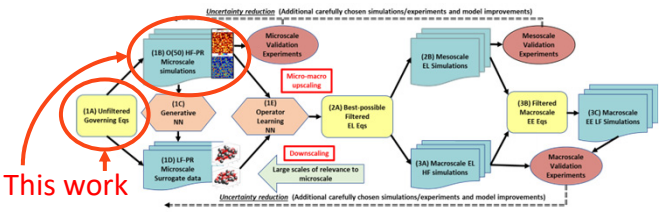

This work

Less Resolution Requires More Closure

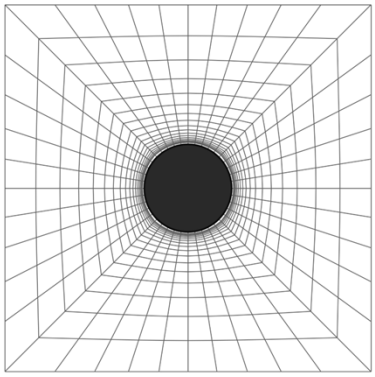

Particle-resolved

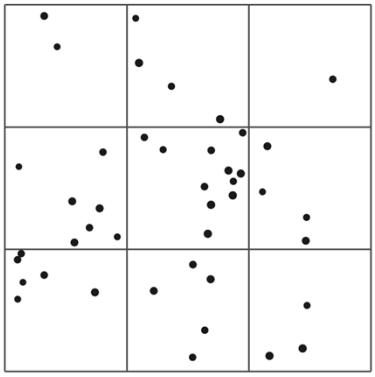

Euler-Lagrange

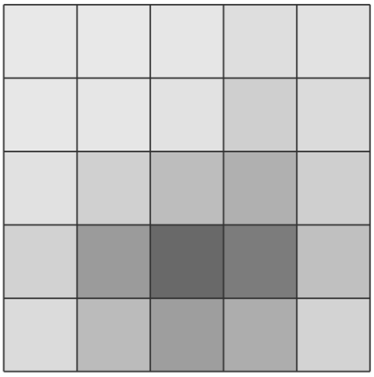

Euler-Euler

CM³C

Less Resolution Requires More Closure

MEAN

What is the average response conditioned on shared descriptors such as Reynolds number, volume fraction, or temperature?

STOCHASTIC

What distribution of plausible responses is consistent with those descriptors and a temporal model?

DETERMINISTIC

What response belongs to this particle after conditioning on its actual neighbors, direction, size, and state?

The distinction is the information supplied to the model, not simply the regression method.

CM³C

Tandem Droplets Expose Realized Front-rear Coupling

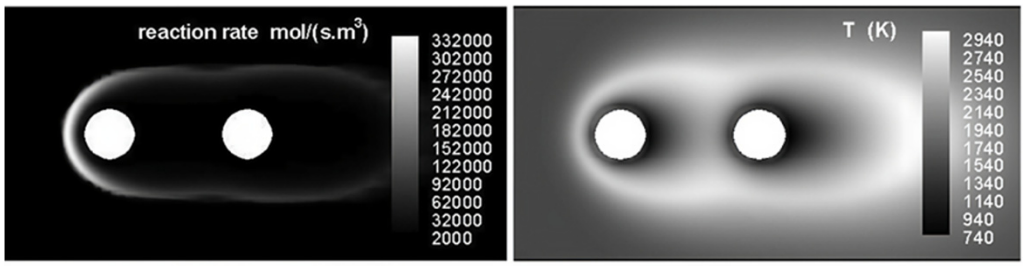

(a) Reaction rate for case 1 (b) T for case 1

Reaction-rate and temperature fields for tandem droplets in crossflow. Sirignano, Case 1.

UPSTREAM DROPLET

Creates the wake and changes velocity, temperature, oxidizer, and fuel vapor before they reach the second droplet.

DOWNSTREAM DROPLET

Experiences lower oxidizer access and a modified reaction rate because it sits in the realized wake.

MODELING POINT

A mean or unconditioned random model does not know which droplet is upstream unless the geometry is supplied.

CM³C

Deterministic Modeling Through Flow Prediction

Predict the local fields, then derive the quantities of interest

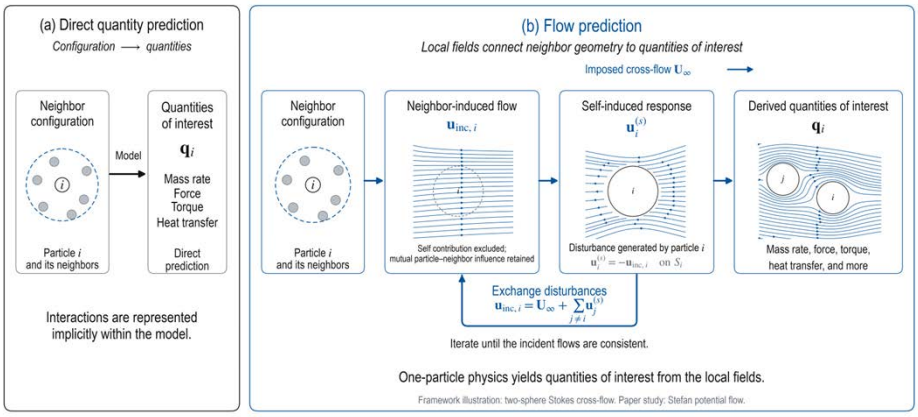

A deterministic model resolves how neighbors change the local fields

CM³C

Example Problem: Thermal Advection-Diffusion

- Reference particle is deeply buried in a large cloud
- Isolate inner boundary layer response via boundary conditions

Study Advection-Diffusion equation:

$$\nabla^2\phi = 0, \qquad \mathbf{u} = -\nabla\phi,$$
$$\nabla^2\theta + \mathrm{Pe}\nabla\phi\cdot\nabla\theta = 0$$

Boundary conditions:

$$(\phi,\theta)_{\mathrm{RP}} = 1,$$
$$(\phi,\theta)_{\mathrm{neighbors}} = (\phi,\theta)_{\infty} = 0$$

Temperature mapping:

$$\theta = \left(1 - e^{-Pe_T\phi}\right)/\left(1 - e^{-Pe_T}\right)$$

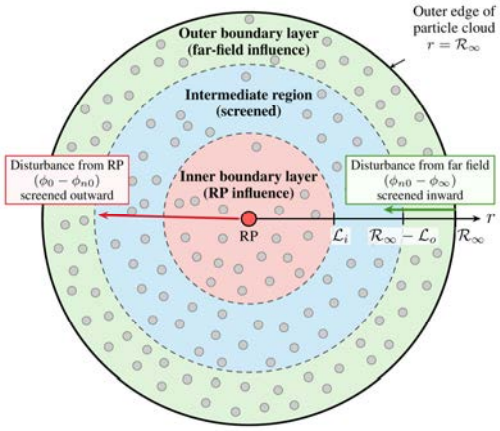

3D RESULTS WITH A 2D CONTOUR; ALL DROPLETS HAVE SAME RADIUS

CM³C

Strength of Deterministic Models

- Neighbor positions and states affect the reference droplet's required heating
- 256 equal-size droplets, 3% volume fraction, 10,000 random configurations (realizations)
- Thermal problem at Pe = 3, γ = 0.25, B = 3; first order in fluctuation amplitude

Model prediction (required heating of the reference droplet)

Mean model

One heating value for all configurations

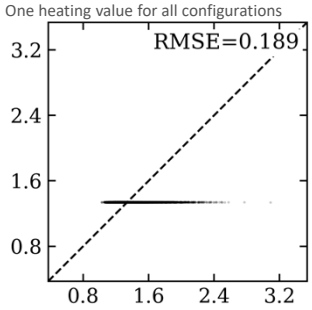

Stochastic model

Mean + matched random spread

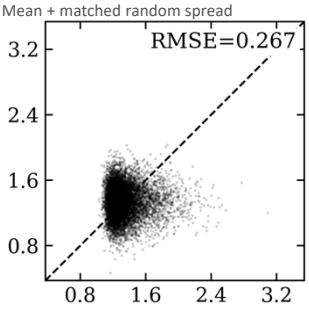

Deterministic model

5 nearest neighbors + mean field

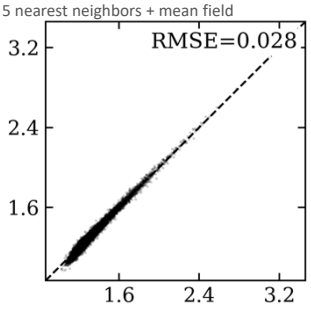

Fully coupled solution (required heating of the reference droplet)

Each point is one configuration. The dashed line denotes exact agreement.

Five neighbors recover the configuration-specific heating requirement

CM³C

Role of The Nearest Neighbors

- A few neighbors recover much of the configuration-specific heating response
- Same 256-droplet, 3% ensemble; 10,000 configurations with fluctuating surface states
- Thermal problem at Pe = 3, γ = 0.25, B = 3; first order in fluctuation amplitude

Model prediction (required heating of the reference droplet)

1 neighbor **5 neighbors** **20 neighbors**

Resolve only the listed nearest neighbors (hard truncation).

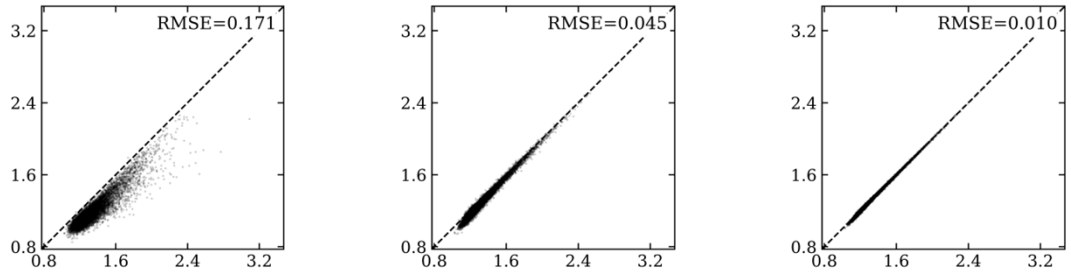

Fully coupled solution (required heating of the reference droplet)

Additional neighbors reduce the remaining particle-to-particle prediction error.

More neighbors recover the remaining variation in required heating

Summary and Future Work

Summary

- Potential flow provides a controlled test of deterministic coupling.
- Nearest neighbors recover the strongest configuration-specific features.

Future work

- Advection–diffusion including heat source terms; Shvab–Zeldovich formulations for combustion (thin flame)
- Finite-Reynolds-number flow with coupled transport
- Reacting droplet clouds with resolved flames

Goal: neighbor-aware models for group combustion

CM³C

Do you have any questions?



NNSA
National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY



The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

N Particle Flow-Prediction: Data Driven Self-Consistent Machine Learning

Raul Valle, Erica Lindbeck, Joel Harley

Electrical and Computer Engineering Department

University of Florida

UF UNIVERSITY of FLORIDA · PURDUE UNIVERSITY · The University of Texas at Austin Oden Institute for Computational Engineering and Sciences

NNSA National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

Our piece of CM³C: a learned stand-in for particle-resolved flow

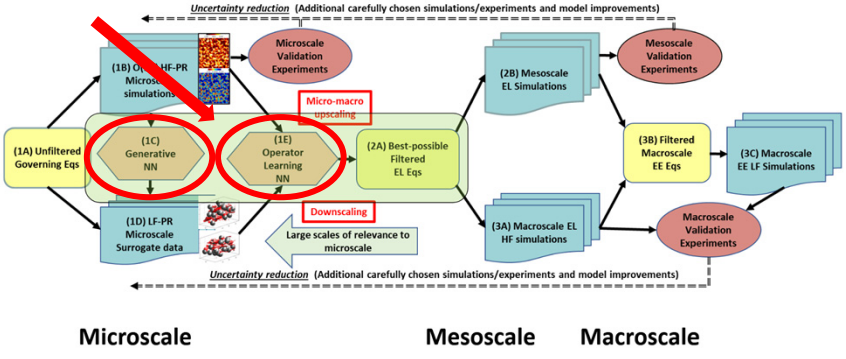

PR-DNS (1B)
particle-resolved simulations: the reference data

Training data
fluid velocity at points between particles

E(3)-equivariant surrogate (1C/1E)
one network, shared by every particle

Fast flow prediction
flow at ~40k points in under a minute on one GPU

Closures for Euler–Lagrange (2A)
proposed use; not yet demonstrated

Since starting: GPU training at scale (HiPerGator, LLNL), a damped alternating schedule for the latent state update, and exhaustive evaluations across simulations.

UF UNIVERSITY of FLORIDA

PR-DNS resolves every particle: 1,000 particles ≈ 14,000 core-hours

And Case 13 is one of 48 simulations in the database

64 to 1,000 particles · volume fraction 1.7–26% · particle Re ≈ 10–67 · density ratio 2.56–256

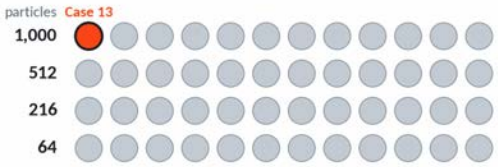

each run: about 1,800 to 29,000 CPU core-hours
7 hours to 8.4 days on 80–400 cores
10,000 to 140,000 time steps

all 48 together: more than 400,000 CPU core-hours (over 45 core-years)

per-run costs: the 36 cases with one run log (solver time without I/O); the total is a lower bound

PR-DNS (C. Lott, NekIBM) · Case 13, snapshot t = 7 · solver nodes, no interpolation · costs from the NekIBM run logs

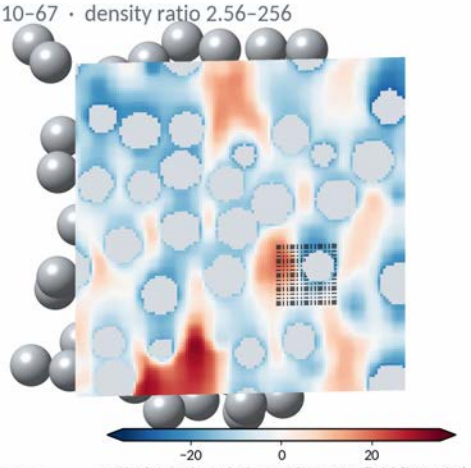

CM³C

UF UNIVERSITY of FLORIDA

Data pre-processing: one PR-DNS snapshot → training examples

Data pre-processing for the model

done once, before any training · PR-DNS snapshot, Case 13 (t = 7)

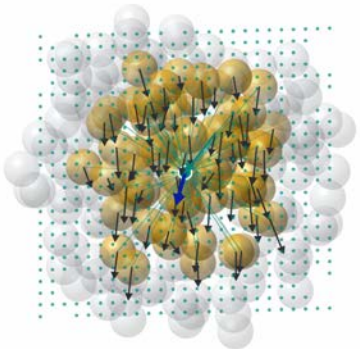

1. **a point x in the fluid**
2. **its 64 nearest particles**
   up to 3.2 d away
3. **inputs: where they are, how they move**
   x − r_j , q_j (speed, direction) , ρp/ρf , φ_j
4. **target: the PR-DNS velocity at x**
   u(x) − ⟨u_f⟩
5. **repeat at 39,875 points**
   every 5th solver node: a d/3 lattice in the fluid
6. **for each of the 15 snapshots**



9 train 3 development 3 test, locked

→ **the training data for the model**

Case 13 · PR-DNS (C. Lott) · development snapshot t = 7 · lattice and targets as built by np_train.build_snapshot

CM³C

Refinement: train the network, then update what each particle sees

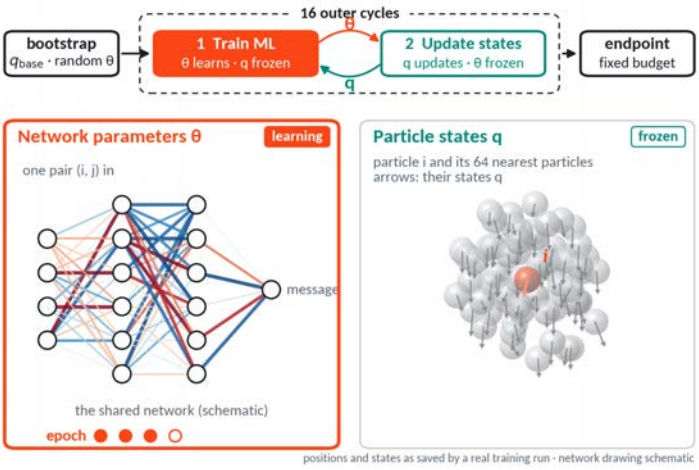

What the state is

Each particle carries a state q. It starts as its velocity relative to the mean fluid, v − ⟨u_f⟩; a latent representation is learned.

One cycle

1. Freeze the states; train the network 4 epochs.

2. Freeze the network; update every state once:

$$q \leftarrow (1-\alpha)\, q + \alpha\, (q_{base} - W_\theta(q)),\ \ \alpha = 1/8$$

* At inference the network is frozen and only step 2 is iterated.

CM³C

Refinement: train the network, then update what each particle sees

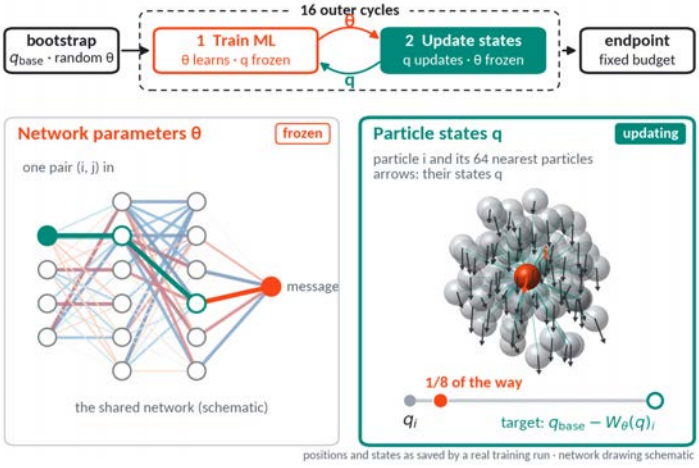

What the state is

Each particle carries a state q. It starts as its velocity relative to the mean fluid, v − ⟨u_f⟩; a latent representation is learned.

One cycle

1. Freeze the states; train the network 4 epochs.

2. Freeze the network; update every state once:

$$q \leftarrow (1-\alpha)\, q + \alpha\, (q_{base} - W_\theta(q)),\ \ \alpha = 1/8$$

* At inference the network is frozen and only step 2 is iterated.

CM³C

One shared equivariant network; per-particle messages add up

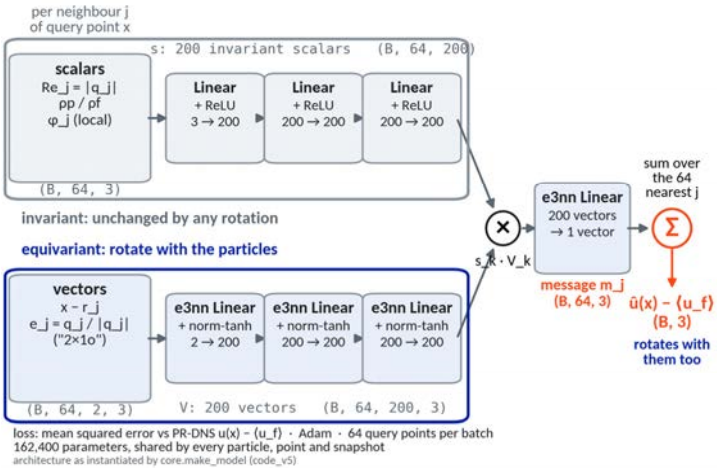

Rotate the particles and their velocities…

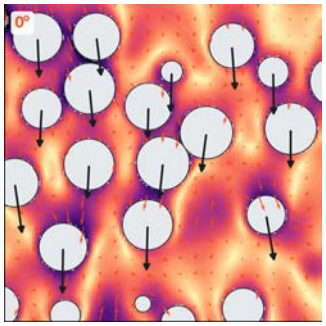

black: particle velocities · orange: predicted flow

…and the predicted flow rotates with them, exactly

max difference 6×10^{-4} on a field of up to 29 (float32 round-off)

Why rotational symmetry matters

Equivariant by construction: exact at any angle
Built in, not learned: no data spent on orientations
No built-in vertical: gravity enters only through the velocities

CM³C

Training alternates: fit the network, then update particle states

What if the particles never update what they see?

Same network, same data, same training budget. One switch: the state update.

state update OFF

CM³C

The winning refinement recipe: 4 epochs, 1 update, 16 cycles

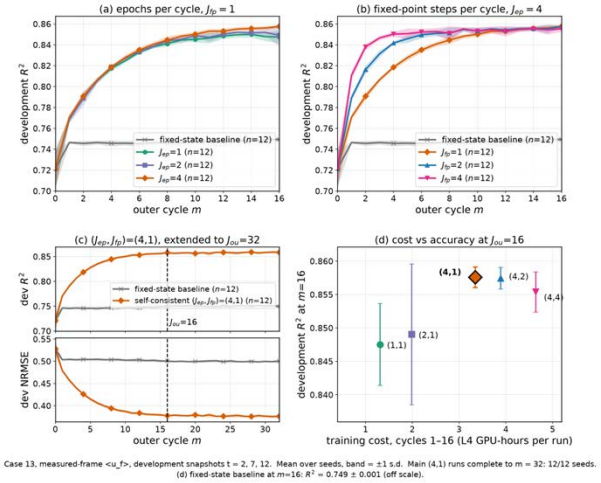

Case 13, measured-frame <u_f>, development snapshots t = 2, 7, 12. Mean over seeds, band = ±1 s.d. Main (4,1) runs complete to m = 32: 12/12 seeds. (d) fixed-state baseline at m=16: R^2 = 0.749 ± 0.001 (off scale).

4
training epochs per cycle
tested 1, 2, 4: most accurate

1
state update per cycle
tested 1, 2, 4: same end point, least cost

16
cycles
tested up to 32: nothing gained after 16

α = 1/8
damping of each update
undamped (α = 1) does not converge

64
neighbours per point
up to 128: slightly better, more cost

50
surface points per particle
14 or 26 equal; centre only −0.008 R^2

All chosen on development snapshots (12 seeds each); test snapshots stayed locked. Left: the evidence (manuscript Fig. 4).

CM³C

Locked test: the update lowers the error everywhere, most between particles

One particle's learned response

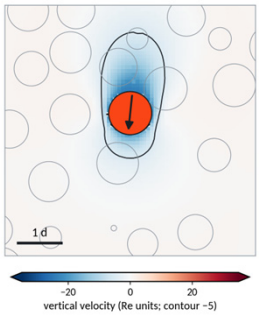

Σ over the 64 nearest particles, at every fluid point, gives the reconstruction →

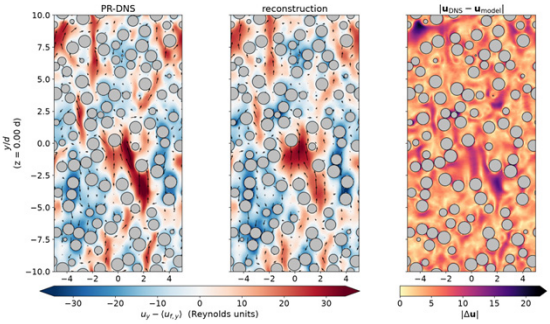

Case 13, locked test snapshot t = 9 (pre-registered evaluation), seed 7057, cycle 16. Vertical velocity relative to the mean fluid on the plane z = 0 (x/d horizontal, y/d vertical); right panel, the error magnitude |Δu|.

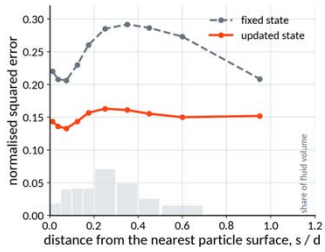

All 5.07 M fluid grid nodes of this snapshot: normalized error is about 41% lower overall due to refinement; the largest drop is 0.3–0.7 d from the surfaces, in the gaps between particles.

Every prediction here is a composition: one learned particle response, summed over each point's 64 nearest particles. With each particle's state updated, that composition lowers the locked-test error at every distance from the particles.

Left: the same network (seed 7057, cycle 16), development snapshot t = 7; the colour scale matches the reconstruction.

CM³C

Updating particle states improves held-out accuracy: R^2 0.75 → 0.85

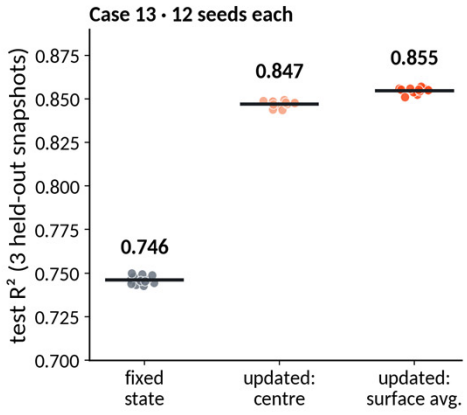

Test : fixed state is outperformed consistently by both refinement types.

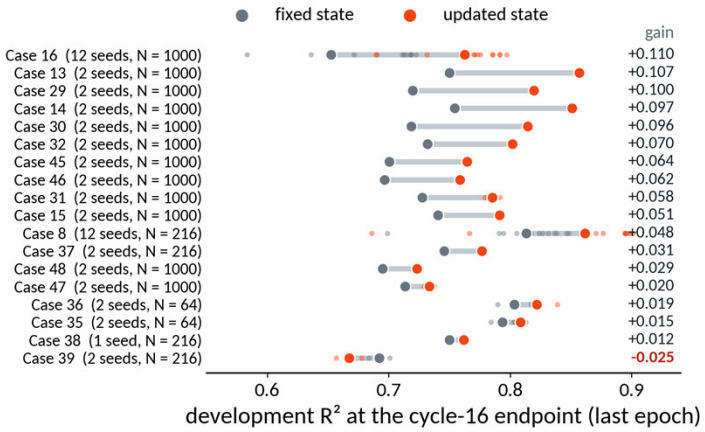

Development snapshots (18 cases so far): the update helps in 17 of 18 at the last epoch; Case 39's development R^2 swings dramatically between epochs

CM³C

Inference is cheap, but no end-to-end computational gain is shown yet

PR-DNS, Case 13 (measured)

61 h	≈ 4.1 h
the whole run	**one snapshot interval**
224 CPU cores	224 CPU cores

Our model, Case 13 (measured)

≈ 3 h	≈ 41 s
training, once	**inference, one snapshot**
1 GPU	1 GPU

Not a saving yet

inputs and reconstructed snapshots still come from PR-DNS

It pays off when

it is reused on configurations PR-DNS has not run, with inputs from Euler–Lagrange (not shown yet)

Not optimised yet, and already much cheaper to run

Already cut: a training epoch from 12–13 min (earlier CPU code) to 17–18 s on one GPU, ≈ 40×

Already cut: 16 cycles instead of 32, half the training

Next: a spatial index for the neighbour search · 14 surface points instead of 50 · lighter GPU memory

CM³C

UF UNIVERSITY of FLORIDA

Takeaway

We contribute a **long-term component** to the CM³C pipeline

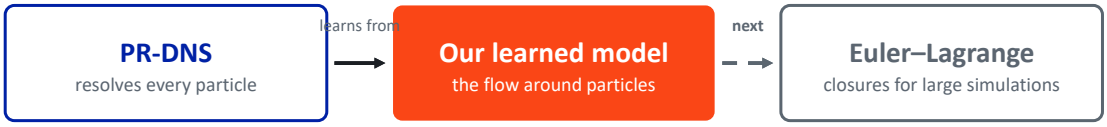

Others can replicate it

one fixed recipe, pre-registered tests

hours of training on one GPU for a new PR-DNS case

Designed to drop in

particle positions and velocities (and the mean fluid velocity) in, **fluid velocity anywhere** out

under a minute per snapshot on one GPU · **$R^2 \approx 0.85$** on held-out PR-DNS flow

Next: what the learned flow tells us about the physics →

CM³C

CM³C

Questions?

UF UNIVERSITY of FLORIDA · PURDUE UNIVERSITY · The University of Texas at Austin Oden Institute for Computational Engineering and Sciences

Backup: model and training details (as in the code)

Item	Setting (np_train.py code_v5, core.make_model)
Inputs per neighbour	x − r_j (3), e_j (3) [2×1o]; Re_j = |q_j|, ρp/ρf, ϕ_j (local, r = 3 d) [3 scalars]
State	q_base = v_j − ⟨u_f⟩ (Reynolds units); ⟨u_f⟩ = measured mean fluid velocity of the snapshot
Neighbours	64 nearest particle centres to each query point (minimum image, periodic box)
Scalar branch	Linear 3→200, ReLU, Linear 200→200, ReLU, Linear 200→200, ReLU
Vector branch	3 × [e3nn o3.Linear (→ 200×1o) + NormActivation(tanh)]
Combine / output	ElementwiseTensorProduct(200×0e, 200×1o) → o3.Linear 200×1o → 1×1o; sum over 64 (no gate)
Parameters	162,400
Data (Case 13)	15 snapshots; 39,875 fluid lattice points each (every 5th solver node, d/3 spacing)
Split	t mod 5 = 2 development (3), t mod 5 = 4 test (3, locked), else train (9)
Optimisation	MSE; batch 64 points; 6 snapshots per epoch; bootstrap 4 epochs Adam 1e-3
Cycles	16 × [state update α = 1/8, surface mean over 50 Lebedev points; 4 epochs Adam 1e-4, plateau(10, 0.1)]
Selection	pre-registered endpoint cycle 16 (epoch 68); 12 seeds; no early stopping

CM³C

Backup: learning curves and the training schedule

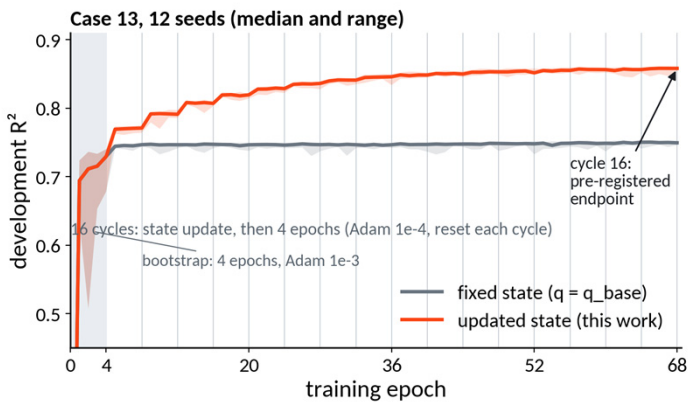

Per epoch: 6 training snapshots × ~40k points, shuffled, batches of 64 points; MSE loss; Adam.

Bootstrap: 4 epochs, lr 1e-3, states fixed at q_base.

16 cycles: one state update q ← (1 − α)q + α(q_base − W_θ(q)), α = 1/8, then 4 epochs with a fresh Adam (lr 1e-4, reduce-on-plateau).

W_θ(q): the flow the neighbours induce, averaged over each particle's surface (50 points).

Selection: fixed, pre-registered endpoint (cycle 16, epoch 68); no early stopping; development snapshots for monitoring; test snapshots locked.

CM³C

Backup: rotation equivariance, checked on the real network

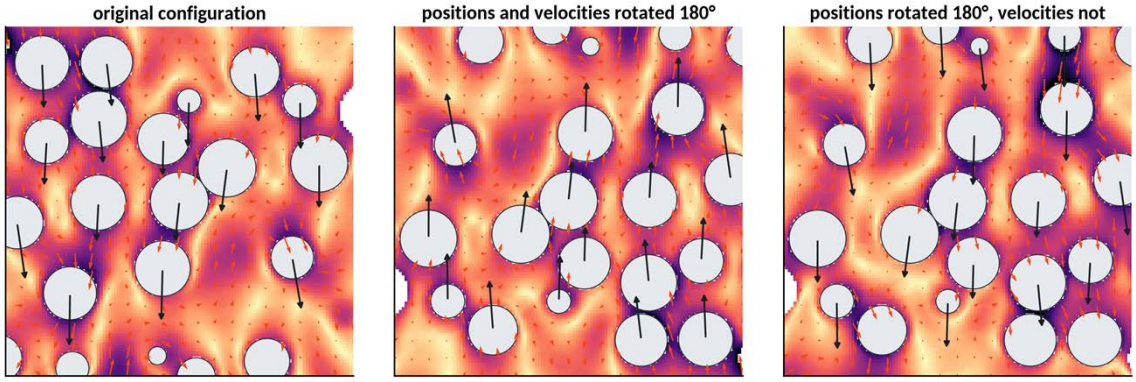

Real network (Case 13, seed 7057, cycle 16) on the 392 particles within 6 d of particle 178 (dev t = 7). Colour: in-plane speed of the predicted flow; arrows: in-plane flow (orange) and particle states (black). Max |F(R·inputs) − R·F(inputs)| = 5.8e-04 over 3 test angles.

Rotating positions **and** velocities rotates the predicted flow exactly (middle = left turned by 180°). Rotating positions alone does not (right): the particle velocities carry the direction — there is no built-in gravity axis.

Backup: where each cost number comes from

PR-DNS, Case 13 — NekIBM run log: "Running on 14 nodes with 224 tasks"; total solver time 2.198e5 s (61.1 h); 30,000 steps; 7.3 s per step on average (7.1 s near steps 26,000–28,000); 224 × 61.1 h = 13,700 core-h. Snapshots used for training are stored every 2,000 steps (times 2, 4, …, 30), ≈ 4 h on 224 cores each. Hardware of the PR-DNS run not recorded here.

Training — DONE.json of each run: wall-clock and torch.cuda.max_memory_allocated on one NVIDIA L4 (HiPerGator hpg-turin). Case 13: 3.0–3.3 h per run (16 cycles); 46-case study runs: 3–8 h, 3.6–10.2 GB.

All CFD experiments since 1 Sep 2026 — Slurm accounting on HiPerGator (u1e3nn jobs only): 9,285 jobs, ≈ 5,290 GPU-h (L4), ≈ 22,300 allocated CPU core-h. LLNL (Tuolumne) usage not included.

Prediction — timing job on 29 Oct (one L4, Case 13 dev snapshot t = 7, Model_16, median of 5): field at 39,875 points 3.74 s (peak 3.0 GB); surface state update 4.66 s; centre update 0.067 s; 500,000 random fluid points 46.4 s → ≈ 7.8 min for 5.07 M (linear extrapolation). Earlier study (R71, 4 L4 hosts): ≈ 1.0 µs per network evaluation, state-update cost ∝ N^1.006 and M^1.09; the brute-force neighbour search grows with N.

Break-even N = (C_data + C_train) / (C_PR-DNS − C_ML) is only meaningful for a comparable workload (same output, same inputs available). Needed measurement: an Euler–Lagrange run with the learned closure vs PR-DNS for the same configuration, cost to equal accuracy.

CM³C

CM³C: Center for Multiscale Modeling of Multiphase Combustion

T7. Plans for the Upcoming Year and Beyond

Technical Manager: Thomas L. Jackson

Department of Mechanical and Aerospace Engineering

UF UNIVERSITY of FLORIDA PURDUE UNIVERSITY The University of Texas at Austin Oden Institute for Computational Engineering and Sciences

UF UNIVERSITY of FLORIDA

Questions for Y2 and Beyond

Q1. Referring back to the simulation roadmap and risks presented yesterday, what is your **year 2 predictive simulation**? If it has **changed** based on what you have learned in your first year, describe how and why.

Q2. What **major activities are required**, and again, if these plans have evolved, how and why have they done so?

Q3. What is the **biggest risk**, i.e., what keeps you up at night?

Q4. What was your **most pleasant surprise**, i.e., what lets you sleep well at night?

Q5. What 3 **takeaway messages** would you like the review team to leave with

Integration of Physics and Simulations

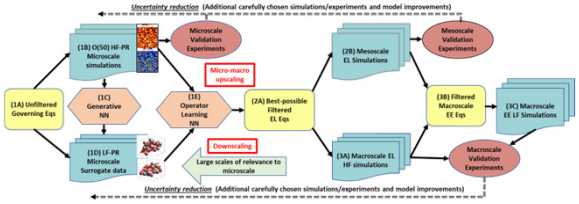

Integration Roadmap

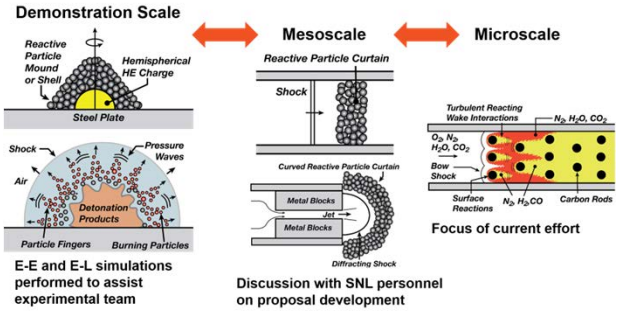

Physics & Simulations

CM³C

Q1&2. Demonstration Problem – Hero Runs

- Ingredients
  1. High explosive (HE) hemispherical charge surrounded by reactive carbon (C) or magnesium (Mg) particles
     a. Molten Mg burns within vapor diffusion flames
     b. Carbon undergoes heterogeneous surface and gas-phase reactions
  2. 12mm and 30mm diameter hemispherical PETN HE
  3. Particle bed thickness 1mm to 10mm
  4. Particle volume fractions 50% and higher
  5. Domain large enough to capture fingering instabilities
- Two tracks
  1. Particle Resolved (immersed boundaries)
  2. Euler-Lagrange (point particles)
- Y5 – push both tracks to computing limits

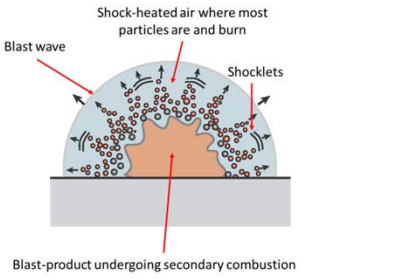

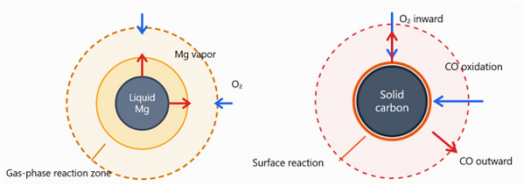

CM³C

Q1&2. Key Physics – Y1 Accomplishments

- Several major objectives satisfied in Y1
  1. EOS
  2. Heterogeneous surface reactions
  3. Euler-Lagrange framework
  4. Modest Y1 Hero simulations
  5. Euler-Euler – out-of-date
  6. Exascale plan
     - b. Baseline code scales to all of Frontier and Tuolumne
     - c. Weak scaling performed for every major new physics module

	Y1	Y2	Y3	Y4	Y5
Gas-phase physics and MFC development					
Gas phase	LES	LES	LES	LES	LES
EOS	Ideal	JWL	BKW	BKW	BKW
Closure model	Standard Smag.	Dynamic Smag.	Dynamic Smag.	ML-based	ML-based
Chem. reactions	$\mathcal{O}(1)$	$\mathcal{O}(10)$	$\mathcal{O}(10)$	$\mathcal{O}(100)$	$\mathcal{O}(100)$
Wall effects	Inviscid	Inviscid	Inviscid	Inviscid	Viscous
Grid size	10^{12}	10^{12}	10^{12}	10^{12}	10^{13}
Euler–Lagrange MFC developments and details					
Dispersion (EL)	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg
Concentration	Dilute	Dilute	Dense	Dense	Dense
Contact model	None	Soft sphere	Frictional soft sphere	Frictional soft sphere	Frictional soft sphere
Force model	Standard	Standard	ML-based	ML-based	ML-based
Particles	10^{10}	10^{11}	10^{12}	10^{13}	10^{13}
Load balancing	Simple	Bin-based	Bin-based	Trace-based	Trace-based
Hero Sims (EL)	2	2	4	4	4
Euler–Euler MFC developments and details					
Dispersion (EE)	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg	Particles: C Droplets: Mg
Distribution	Monodisperse	Polydisperse	Polydisperse	Polydisperse	Polydisperse
Grid size	10^{10}	10^{10}	10^{10}	10^{10}	10^{10}
Grid refinement	Static	Static	AMR	AMR	AMR
Concentration	Dilute-dense	Dilute-dense	Dilute-dense	Dilute-dense	Dilute-dense
Contact model	Granular	Granular	Granular	ML-based	ML-based
Force model	Standard	Standard	ML-based	ML-based	ML-based
EE Sims.	50	100	300	600	1000
Exascale plan					
Architecture	GPU (Lassen) APU (El Capitan)	APU (El Capitan)	APU (El Capitan)	APU (El Capitan)	APU (El Capitan)
Machine utilization	~100%	~100%	~100%	~100%	~100%

CM³C

Q1&2. Key Physics – Y1 Simulation Highlights

- Particle-Resolved simulation
  1. 100,000 magnesium particles; 45-micron diameter; 20 pts per diameter
  2. 10% initial particle volume fraction
  3. 34 billion grid cells; Δh=2.25 microns
  4. 4096 AMD GPUs for 8 hours on Frontier -> 8.3M cells per GPU; 16k immersed boundaries per GPU
- Euler-Lagrange simulation
  1. 70 million individually tracked Lagrangian particles
  2. 20% initial particle volume fraction
  3. 32 million grid cells; Δh=500 microns
  4. 10 nodes on LLNL Tuolumne, 40 AMD MI300A GPUs -> 0.8M cells per GPU
- Both use high pressure gas as surrogate for HE
- Both use inert particles

Daniel Vickers

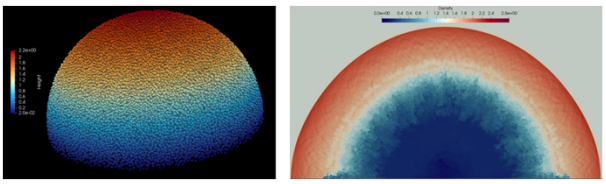

Thierry Daoud

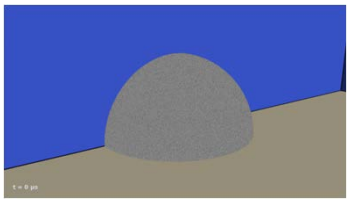

CM³C

Q1&2. Y1 Weak Scaling

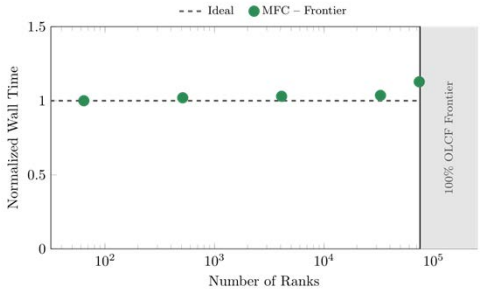

* Caveat: Different simulation campaign, unclear result

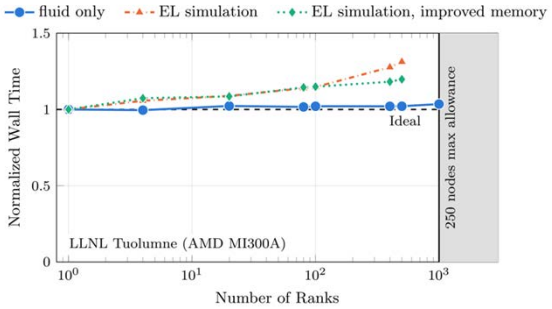

CM³C

Q1&2. Y2 Predictive Simulation and Major Activities Required

- Physics to be included in Y2 for both EL and PR tracks
  1. Program burn for propagating radial detonation wave in HE
  2. Proper mixture rule at the explosive products-air interface
  3. Chemistry for carbon and magnesium particles
- Euler-Lagrange
  1. Complete set of up-to-date force models
  2. Move towards ML models

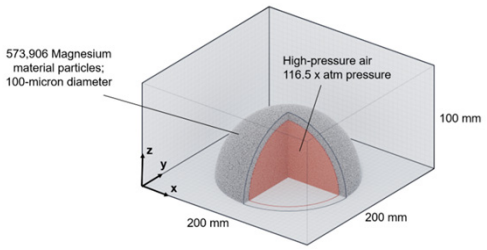

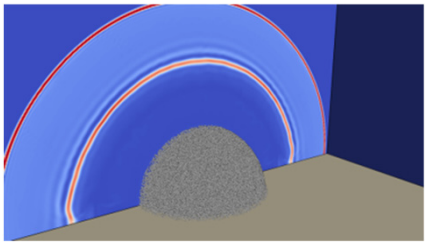

CM³C

Q1&2. Y2 Proposed Hero Simulations

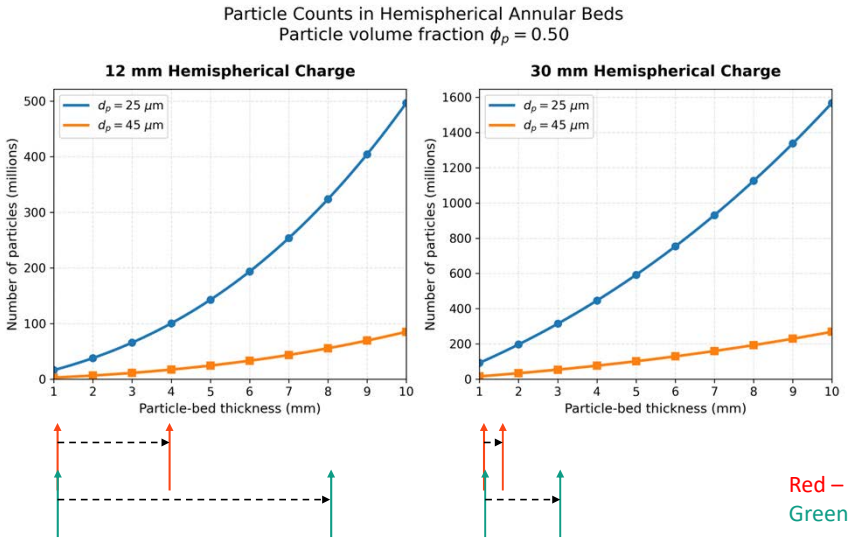

Red – PR
Green - EL

CM³C

Q1&2. Changes Made Based On Lessons Learnt - I

- We will pursue detailed chemistry in PR-DNS
- Particle-resolved IBM simulations will be performed with Millions to Billions of particles (depending on resolution)
- Demonstration experiments/simulations will use a hemispherical packing of particles, instead of free-standing heap
- Radiation is important for carbon rod experiment and we will account for it in the simulations
  1. Its importance for particles will be examined

CM³C

Q1&2. Changes Made Based On Lessons Learnt - II

- Particle temperature effect is more important than what we expected, even in force modeling
  1. Need to include it as a key parameter
- Generative NN ⇒ Flow-Prediction NN
  1. A self-consistent one-particle framework has replaced our CNN approach
  2. We do NOT need to simulate and store surrogate data as originally planned
- Operator Learning NN ⇒ Force-Prediction NN
  1. We will predict all the closure quantities

Q3. Challenges Resolved and Challenges Ahead

- Reduced order chemistry has too much uncertainty
  - MFC can do handled detailed chemistry
- Radiation between particles (or rods) is non negligible
  - Yes, simple radiation will need to be implemented
- Load-balanced exascale EL implementation in MFC
  - MFC is very versatile – we do not see this as a major issue
- Not enough PR data for Generative ML model
  - Running big/bundled simulations with MFC is easy
- Initial compaction phase introduces complex unaccounted physics
- Superposition model fails for finite Re combustion
- Generative + Operator Learning models introduce too much uncertainty
- Availability of computing resources

Q4. Greatest Surprise – Interest & Enthusiasm from Students/Postdocs

Internal Kick-Off Meeting

- December 23, 2025
- 2-hr Zoom
- 28 in attendance
- Overview presentation
  1. History of DOE/PSAAP program
  2. Big-picture view
  3. Key science questions
  4. Introduction of Teams/Tasks
  5. Purpose of weekly student-led meetings
  6. Outline of short-term tasks/goals
- Round-robin of introductions by all present

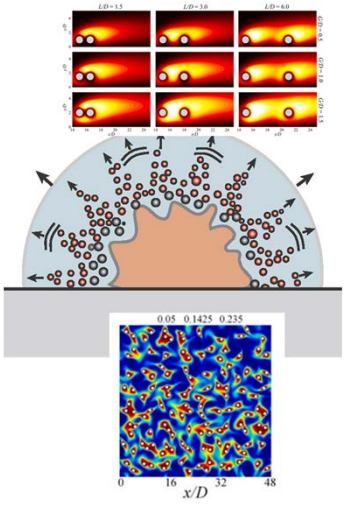

Concept of Group Combustion

Q4. Greatest Surprise – Interest & Enthusiasm from Students/Postdocs

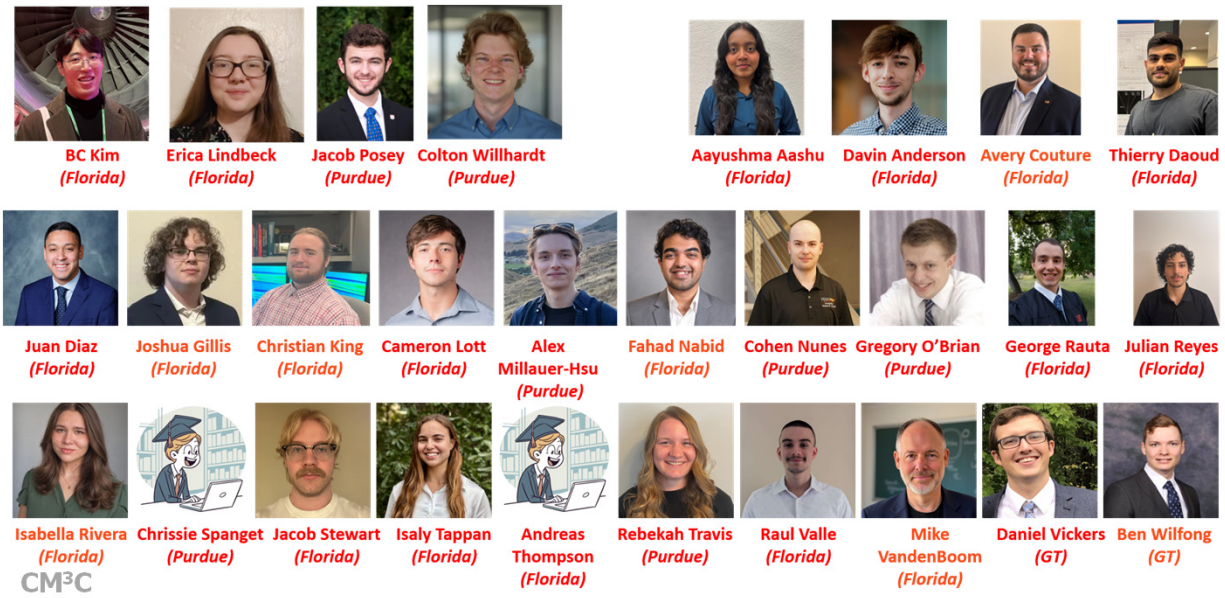

Q4. Greatest Surprise – Ease of Using MFC for Big Simulations

- MFC is well documented and easy to use to get started
  1. Over 4000 case parameters
- Easily runs on Hipergator and Tuolumne
- Numerous example cases
  1. ≈ 200 example cases for 0D – 3D covering wide range of physics
- MFC has a rigorous PR process
  1. Code can be trusted that no one random user has "broken" the code
- Communication between UF and GT/UTA was surprisingly natural
  1. CM3C slack helped facilitate bug reports, ask specific questions, show off latest figure or movie, etc
     a. Response from GT/UTA team is almost instantaneous to questions

Q4. Greatest Surprise – Team Integration!

- Teams
  1. Phys-S: Microscale PRS
  2. Phys-S: Meso-Macroscale EL
  3. CS-Exe: MFC Computational Science
  4. Phys-E: Experiments
  5. ML: Machine Learning
  6. VVUQ: Relay Validation/Workflow
  7. Group Combustion
- All members can easily fit into other teams
- ***Physical distance not a deterrence***

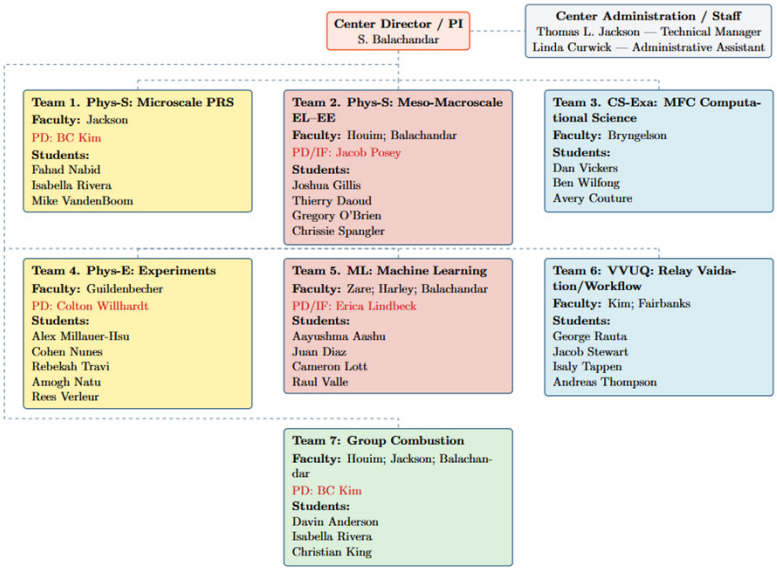

Q4. Greatest Surprise – Team Communication CM³C Slack

- Allows for everyone to be part of the conversation
- 39 users
- Active channels
  1. #general
  2. #gpu-ppiclf
  3. #lc-users (Livermore Users)
  4. #machine-learning
  5. #macro-ee-sims
  6. #macro-el-sims
  7. #macro-exp
  8. #micro-pr-sims
- Frequent crossover with MFC slack

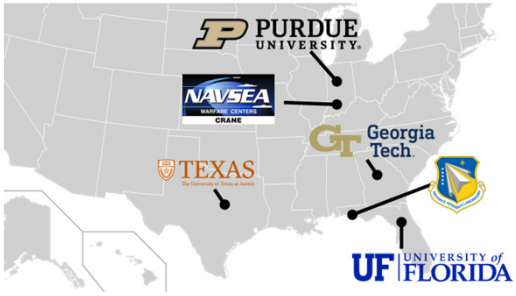

Allows everyone to be part of the conversation

cm3c.slack.com
mflowcode.slack.com

Q5. Takeaway Messages

- What takeaway messages would you like the review team to leave with
  1. Center teams are in place and are functioning well together towards a common goal
  2. Significant code development in Y1, satisfying several Y5 goals
  3. Y2 will see significant database generated for ML for variety of flow configurations
  4. Y2 will see a push towards cutting-edge research into group combustion physics
  5. …

CM³C

Questions?



NNSA
National Nuclear Security Administration

UF UNIVERSITY of FLORIDA

PURDUE UNIVERSITY

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences