CCMT

CENTER FOR COMPRESSIBLE
MULTIPHASE TURBULENCE

TST Review

December 12-13, 2018

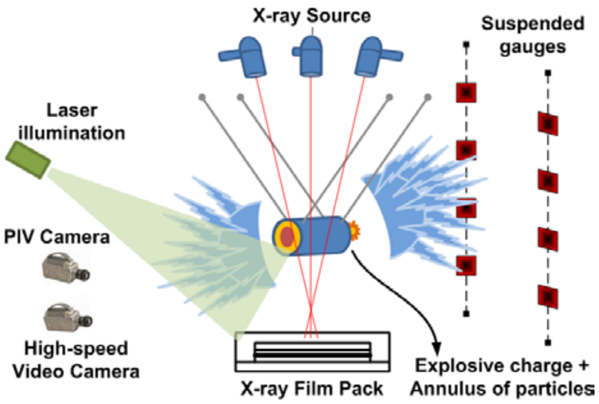

UNIVERSITY of FLORIDA

Center for Compressible Multiphase Turbulence

1180 Center Drive
P.O. Box 116135
Gainesville, FL 32611
Phone: (352)294-2829
Fax: (352) 846-1196

Agenda for TST Site Visit December 12 and 13, 2018

Wednesday December 12, 2018

7:40	Van arriving University Hilton for pickup
7:45	Van leaving University Hilton
8:00-8:45	Full Breakfast (Review team caucus)
8:45-10:00	T1 - CCMT Overview and Integration (Jackson, Bala)
10:00-10:30	T2 – Integrated Simulations (Bertrand Rollin)
10:30-10:45	Coffee break
10:45-11:15	T3 – Experiments (Heather Zunino)
11:15-11:45	T4 – Experiments (Angela Diggs, Kyle Hughes)
11:45-1:00	Lunch
1:00-1:45	T5 – CMT-nek (Jason Hackl, David Zwick)
1:45-2:30	T6 - UB (Chanyoung Park)
2:30-2:45	Coffee break
2:45-3:30	T7 - CS (Tania Banerjee)
3:30-4:15	T8 - Exascale (Herman Lam, Greg Stitt)
4:15-4:20	T9 - Internship presentations by Fred Ouellet
4:20-5:30	Social with Students/Posters (NEB Main Lobby)
5:30	Transportation to University Hilton by CCMT Faculty
6:00-7:30	Dinner at the University Hilton, hosted by CCMT Faculty

Thursday December 13, 2018

7:45	Van pickup at University Hilton
8:00-9:00	Continental Breakfast
9:00-10:45	Student Presentations (15 minutes each for presentations and questions)
	T10 – Mesoscale Simulations (Josh Garno)
	T11 – PIEP Modeling (Chandler Moore)
	T12 – Point-particle vs. Finite-size particle models (Samaun Nili)
	T13 – Load Balancing (Keke Zhai)
	T14 – BE of CMT-nek, Part 1 (Aravind Neelakantan)
	T15 – BE of CMT-nek, Part 2 (Trokon Johnson)
10:45-11:30	Discussions between TST and CCMT PIs
11:30	Box Lunch; Transportation to hotel and/or airport as needed

UNIVERSITY of FLORIDA

CENTER FOR COMPRESSIBLE MULTIPHASE TURBULENCE

TST Review December 12-13, 2018 Attendee List

Faculty

S. Balachandar “Bala”	University of Florida	bala1s@ufl.edu
Rafi Haftka	University of Florida	haftka@ufl.edu
Nam-Ho Kim	University of Florida	nkim@ufl.edu
Herman Lam	University of Florida	hlam@ufl.edu
Sanjay Ranka	University of Florida	ranka@cise.ufl.edu
Greg Stitt	University of Florida	gstitt@ece.ufl.edu
Tom Jackson	University of Florida	tlj@ufl.edu
Siddharth Thakur “ST”	University of Florida	sst@ufl.edu
Bertrand Rollin	Embry-Riddle	rollinb@erau.edu
Ju Zhang	Florida Institute of Technology	jzhang@fit.edu

Review Team

Abhinav Bhatele (Chair)	LLNL	bhatele1@llnl.gov
David Daniel	LANL	ddd@lanl.gov
Maya Gokhale	LLNL	gokhale2@llnl.gov
John Hewson	Sandia	jchewso@sandia.gov
Mark Schraad	LANL	schraad@lanl.gov
Greg Weirs	Sandia	vgweirs@sandia.gov

Others

Andrew Cook	LLNL	cook33@llnl.gov
Angela Diggs	University of Florida	angela.diggs.1@us.af.mil
David Etim	HQ	David.Etim@nnsa.doe.gov
John Feddema	SNL	jtfedde@sandia.gov
Donald Littrell	Eglin Air Force Base	donald.littrell@us.af.mil
Kambiz Salari	LLNL	salari1@llnl.gov

Research Staff

Tania Banerjee	University of Florida	tmishra@cise.ufl.edu
Kei Fujisawa	University of Florida	fujisawa_ocean_eng@yahoo.co.jp
Jason Hackl	University of Florida	jason.hackl@ufl.edu
Nguyen Tri Nguyen	University of Florida	tringuyenttt@gmail.com
Chanyoung Park	University of Florida	cy.park@ufl.edu

Students

Ryan Blanchard	University of Florida	rlb1116@ufl.edu
Sai Chenna	University of Florida	chenna@hcs.ufl.edu
Brad Durant	University of Florida	neoncrash@ufl.edu
Giselle Fernandez	University of Florida	gisellefernandez@ufl.edu
Joshua Garno	University of Florida	jgarno@ufl.edu
Trokon Johnson	University of Florida	johnson@hcs.ufl.edu
Kyle Hughes	University of Florida	kylethughes89@ufl.edu
Rahul Koneru	University of Florida	rahul.koneru@ufl.edu
Tadbhagya Kumar	University of Florida	tkumar12@ufl.edu
Adeesha Malavi	University of Florida	adeeshaw@ufl.edu
Yash Mehta	University of Florida	ymehta@ufl.edu
Chandler Moore	University of Florida	wcmoore@ufl.edu
Aravind Neelakantan	University of Florida	aravindneela@ufl.edu
Samaun Nili	University of Florida	samaunnili@ufl.edu
Frederick Ouellet	University of Florida	f.ouellet@ufl.edu
Carlo Pascoe	University of Florida	poppyc@ufl.edu
Keke Zhai	University of Florida	zhaikeke@ufl.edu
Heather Zunino	Arizona State University	hzunino@asu.edu
David Zwick	University of Florida	dpzwick@ufl.edu

Administration Staff

Hollie Starr	University of Florida	hstarr@ufl.edu

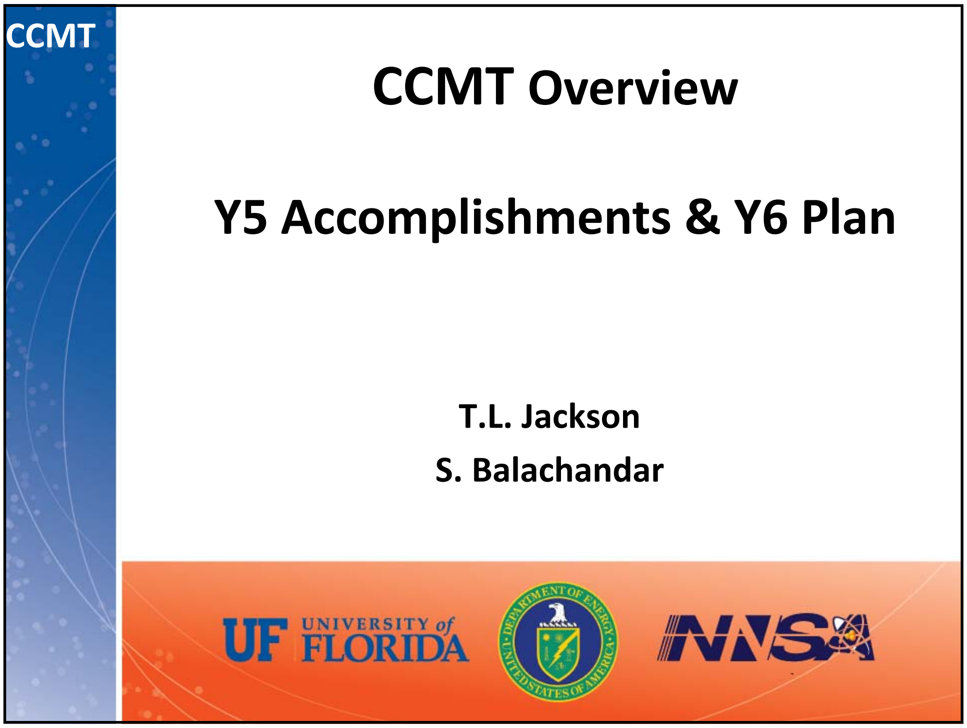

CCMT

CCMT Overview

Y5 Accomplishments & Y6 Plan

T.L. Jackson
S. Balachandar

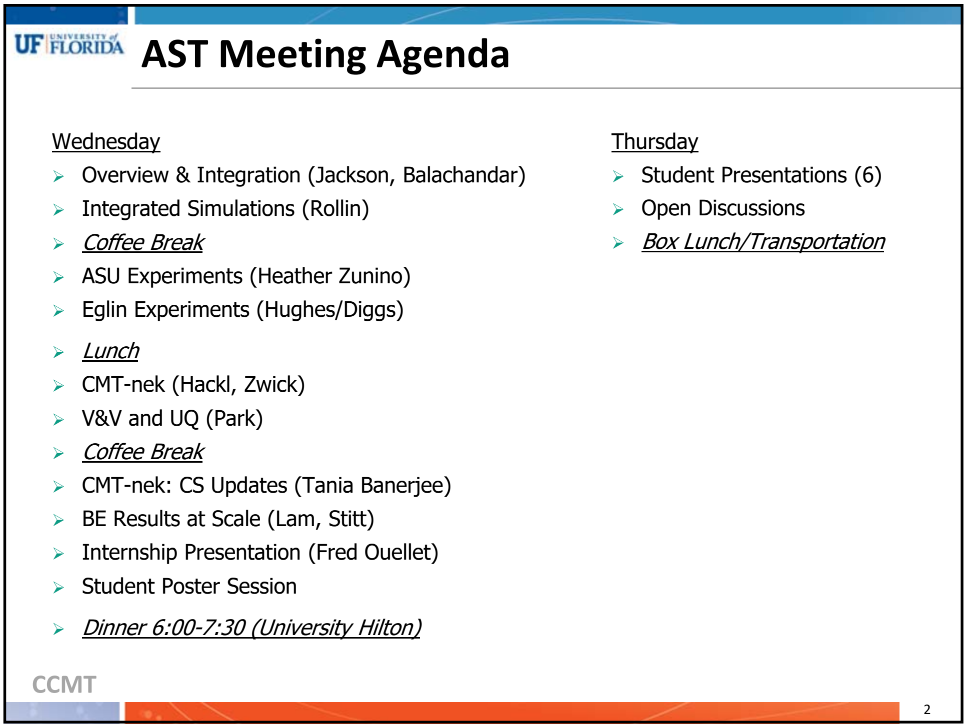

UF UNIVERSITY of FLORIDA

AST Meeting Agenda

Wednesday

- Overview & Integration (Jackson, Balachandar)
- Integrated Simulations (Rollin)
- *Coffee Break*
- ASU Experiments (Heather Zunino)
- Eglin Experiments (Hughes/Diggs)
- *Lunch*
- CMT-nek (Hackl, Zwick)
- V&V and UQ (Park)
- *Coffee Break*
- CMT-nek: CS Updates (Tania Banerjee)
- BE Results at Scale (Lam, Stitt)
- Internship Presentation (Fred Ouellet)
- Student Poster Session
- *Dinner 6:00-7:30 (University Hilton)*

Thursday

- Student Presentations (6)
- Open Discussions
- *Box Lunch/Transportation*

CCMT

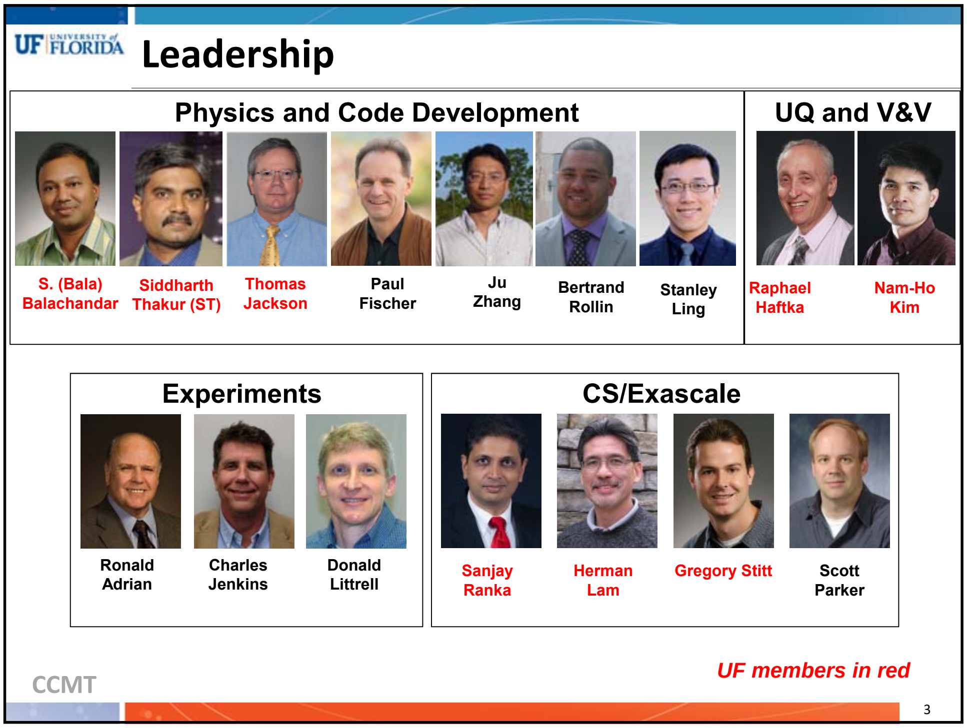

Leadership

Physics and Code Development

S. (Bala) Balachandar
Siddharth Thakur (ST)
Thomas Jackson
Paul Fischer
Ju Zhang
Bertrand Rollin
Stanley Ling

UQ and V&V

Raphael Haftka
Nam-Ho Kim

Experiments

Ronald Adrian
Charles Jenkins
Donald Littrell

CS/Exascale

Sanjay Ranka
Herman Lam
Gregory Stitt
Scott Parker

UF members in red

CCMT

Research Staff & Senior PhD Students

Tania Banerjee
Angela Diggs (Eglin AFB)
Kei Fujisawa
Jason Hackl
Nguyen T. Nguyen
Chanyoung Park
Carlo Pascoe
Sangjune Bae

CCMT

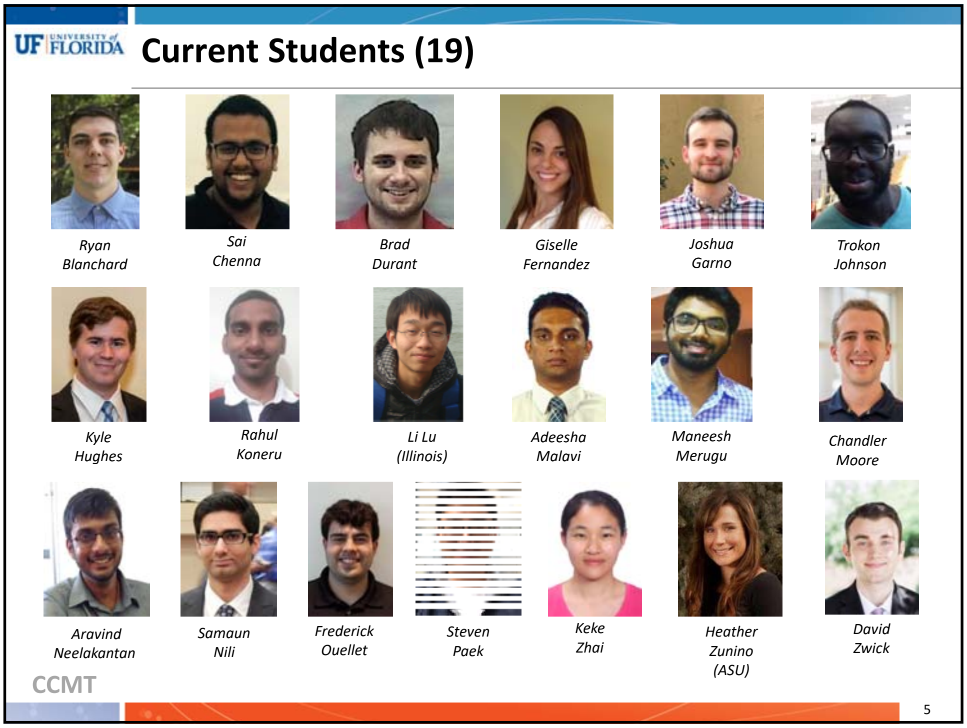

Current Students (19)

Ryan Blanchard
Sai Chenna
Brad Durant
Giselle Fernandez
Joshua Garno
Trokon Johnson
Kyle Hughes
Rahul Koneru
Li Lu (Illinois)
Adeesha Malavi
Maneesh Merugu
Chandler Moore
Aravind Neelakantan
Samaun Nili
Frederick Ouellet
Steven Paek
Keke Zhai
Heather Zunino (ASU)
David Zwick

Internship Program – Completed (19)

Heather Zunino	LANL	May-Aug, 2014	Dr. Kathy Prestridge
Kevin Cheng	LLNL	May-Aug, 2014	Dr. Maya Gokhale
Nalini Kumar	Sandia	March-Aug, 2015	Dr. James Ang
Christopher Hajas	LLNL	May-Aug, 2015	Dr. Maya Gokhale
Christopher Neal	LLNL	June-Aug, 2015	Dr. Kambiz Salari
Carlo Pascoe	LLNL	June-Aug, 2015	Dr. Maya Gokhale
Giselle Fernandez	Sandia	Oct-Dec, 2015	Drs. Gregory Weirs & Vincent Mousseau
Justin Mathew	LANL	May-Aug, 2015	Dr. Nick Hengartner
David Zwick	Sandia	May-Aug, 2016	Drs. John Pott & Kevin Ruggirello

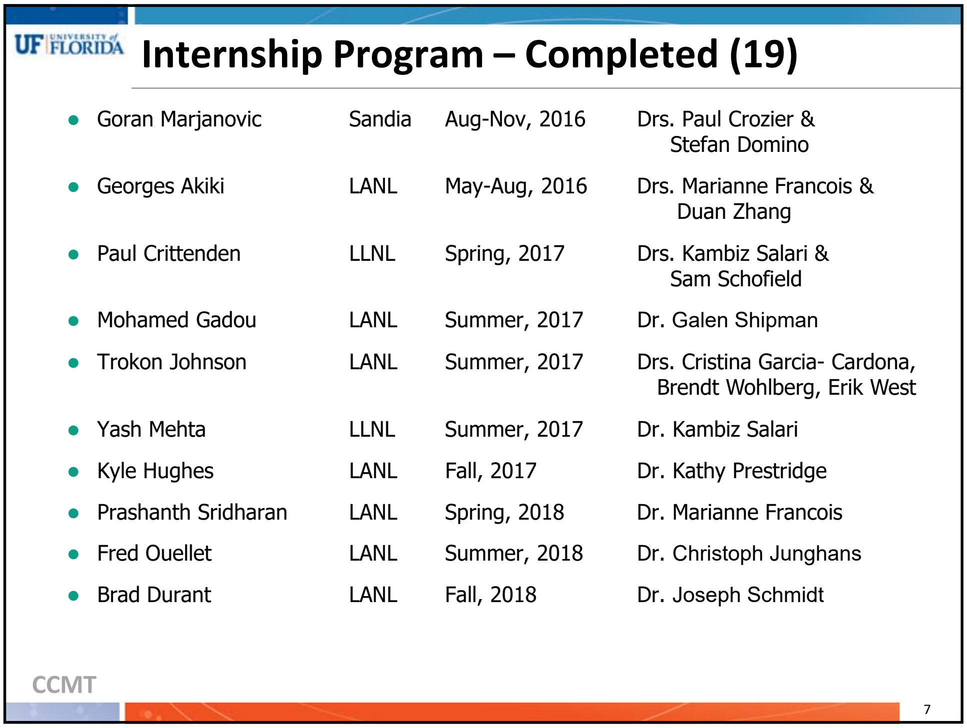

Internship Program – Completed (19)

Goran Marjanovic	Sandia	Aug-Nov, 2016	Drs. Paul Crozier & Stefan Domino
Georges Akiki	LANL	May-Aug, 2016	Drs. Marianne Francois & Duan Zhang
Paul Crittenden	LLNL	Spring, 2017	Drs. Kambiz Salari & Sam Schofield
Mohamed Gadou	LANL	Summer, 2017	Dr. Galen Shipman
Trokon Johnson	LANL	Summer, 2017	Drs. Cristina Garcia- Cardona, Brendt Wohlberg, Erik West
Yash Mehta	LLNL	Summer, 2017	Dr. Kambiz Salari
Kyle Hughes	LANL	Fall, 2017	Dr. Kathy Prestridge
Prashanth Sridharan	LANL	Spring, 2018	Dr. Marianne Francois
Fred Ouellet	LANL	Summer, 2018	Dr. Christoph Junghans
Brad Durant	LANL	Fall, 2018	Dr. Joseph Schmidt

Internship Program – Not Completed

- Joshua Garno, PhD (MAE, Physics and UQ)
- Chandler Moore, PhD (MAE, Physics) (NSF Graduate Research Fellowship)
- Ryan Blanchard, PhD (ECE, BE)

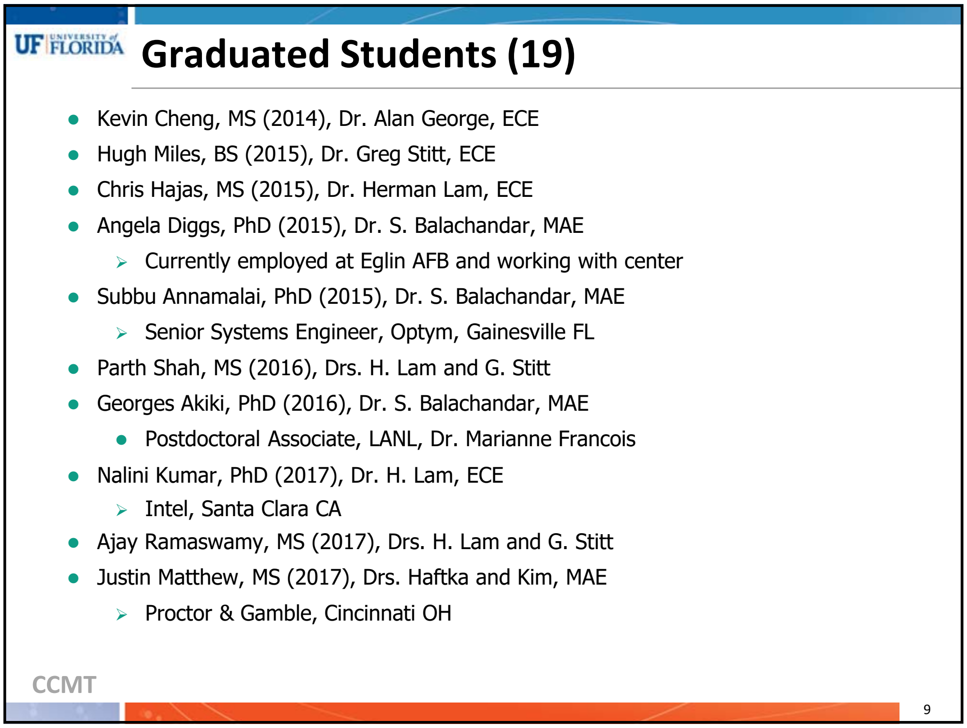

Graduated Students (19)

- Kevin Cheng, MS (2014), Dr. Alan George, ECE
- Hugh Miles, BS (2015), Dr. Greg Stitt, ECE
- Chris Hajas, MS (2015), Dr. Herman Lam, ECE
- Angela Diggs, PhD (2015), Dr. S. Balachandar, MAE
    - Currently employed at Eglin AFB and working with center
- Subbu Annamalai, PhD (2015), Dr. S. Balachandar, MAE
    - Senior Systems Engineer, Optym, Gainesville FL
- Parth Shah, MS (2016), Drs. H. Lam and G. Stitt
- Georges Akiki, PhD (2016), Dr. S. Balachandar, MAE
    - Postdoctoral Associate, LANL, Dr. Marianne Francois
- Nalini Kumar, PhD (2017), Dr. H. Lam, ECE
    - Intel, Santa Clara CA
- Ajay Ramaswamy, MS (2017), Drs. H. Lam and G. Stitt
- Justin Matthew, MS (2017), Drs. Haftka and Kim, MAE
    - Proctor & Gamble, Cincinnati OH

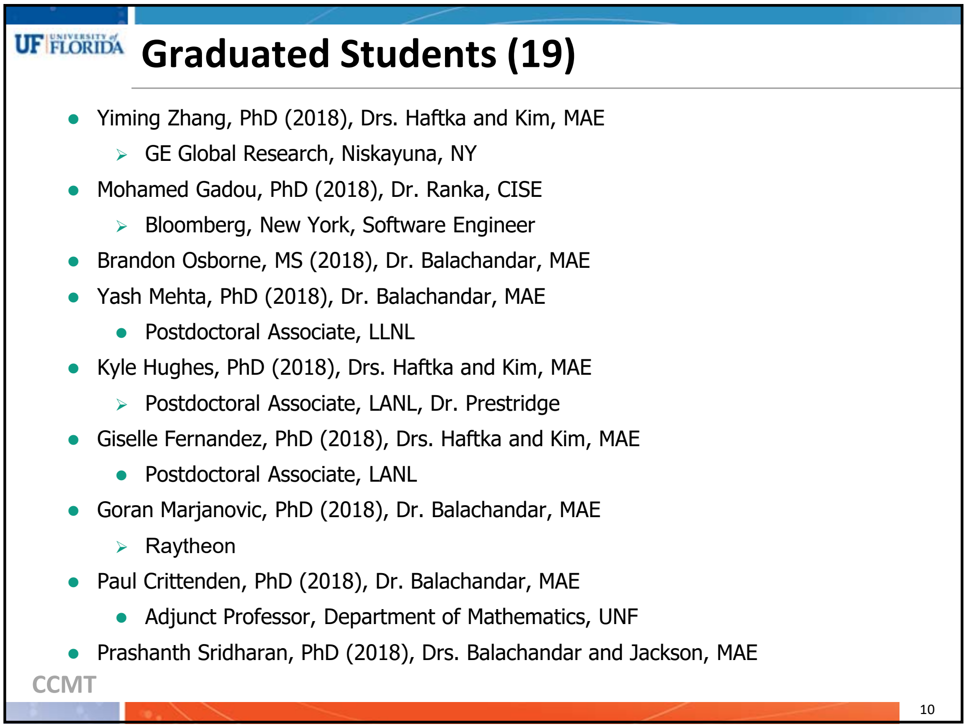

Graduated Students (19)

- Yiming Zhang, PhD (2018), Drs. Haftka and Kim, MAE
    - GE Global Research, Niskayuna, NY
- Mohamed Gadou, PhD (2018), Dr. Ranka, CISE
    - Bloomberg, New York, Software Engineer
- Brandon Osborne, MS (2018), Dr. Balachandar, MAE
- Yash Mehta, PhD (2018), Dr. Balachandar, MAE
    - Postdoctoral Associate, LLNL
- Kyle Hughes, PhD (2018), Drs. Haftka and Kim, MAE
    - Postdoctoral Associate, LANL, Dr. Prestridge
- Giselle Fernandez, PhD (2018), Drs. Haftka and Kim, MAE
    - Postdoctoral Associate, LANL
- Goran Marjanovic, PhD (2018), Dr. Balachandar, MAE
    - Raytheon
- Paul Crittenden, PhD (2018), Dr. Balachandar, MAE
    - Adjunct Professor, Department of Mathematics, UNF
- Prashanth Sridharan, PhD (2018), Drs. Balachandar and Jackson, MAE

PhD Students Expected Graduation Date

- 2019
  - Carlo Pascoe
  - Heather Zunino
  - Rahul Koneru
  - Samaun Nili
  - Fred Ouellet
  - Keke Zhai
  - David Zwick
- 2020+
  - Brad Durant
  - Josh Garno
  - Sai Chenna
  - Trokon Johnson
  - Chandler Moore
  - Aravind Neelakantan
  - Ryan Blanchard

Students & Postdocs Placed at DOE/DOD Laboratories

- Angela Diggs – Eglin AFB (2015)
- Georges Akiki – Postdoctoral Associate, LANL, Dr. Marianne Francois (2016)
- Cameron Stewart – Naval Surface Warfare Center, Indian Head, MD (2018)
- Kyle Hughes – Postdoctoral Associate, LANL, Dr. Prestridge (2018)
- Goran Marjanovic – Raytheon (2018)
- David Zwick – Multiphysics Code Development Computer Scientist, Computational Multiphysics Dept., Sandia National Laboratories (2019)
- Jason Hackl – Naval Surface Warfare Center, Indian Head, MD (2019)
- Giselle Fernandez – Sandia or LANL (2019)
- Yash Mehta – LANL (?) (2019)
- Waiting in line: 3 domestic, 4 international

- Academia: 3 (Georges Akiki; Bertrand Rollin; Paul Crittenden)
- Industry: 3 postdocs (Mrugesh Shringarpure, Subbu Annamalai; Chanyoung Park)
  11 students (3 PhD; 7 BS and MS)

Deep Dive Workshops

- Exascale & CS Issues
  - Feb 3-4, 2015
  - University of Florida
- Multiphase Physics – I
  - Oct 13-14, 2016
  - Tampa FL
- CMT-nek/nek5000 User & Developer
  - April 17-18, 2018
  - Tampa FL
- Multiphase Physics – II
  - Stanford; Florida co-lead
  - October 22-23, 2018
- IUTAM Symposium on Dynamics and Stability of Fluid Interfaces
  - April 2-5, 2018
  - Gainesville, FL
- Workshop on Multiphase Flows
  - November 15-16, 2018
- Trilab Multiphase Workshop
  - 3 days; 2019; initial planning stages
  - S. Balachandar and Duan Zhang (LANL)

CMT-nek/nek5000 User& Developer Meeting

April17-18, 2018

Internal Workshops

- “Dakota - Tutorial” – organized by Chanyoung Park, February 19, 2015
- "Good Software Engineering Practices and Beyond" – organized by Bertrand Rollin, Feb 19, 2015
- “A Boot Camp on CMT-nek” – organized by B. Rollin and J. Hackl, November 29, 2017
- CMT-nek Multiphase Flows – early 2019

CMT-nek Work Camp, 2017

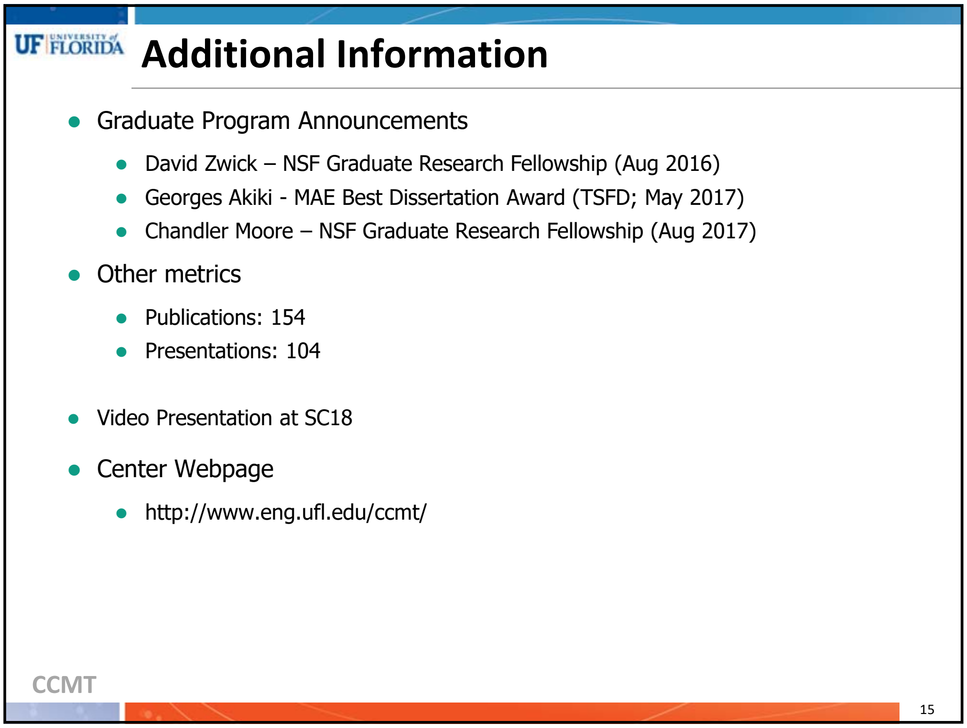

Additional Information

- Graduate Program Announcements
  - David Zwick – NSF Graduate Research Fellowship (Aug 2016)
  - Georges Akiki - MAE Best Dissertation Award (TSFD; May 2017)
  - Chandler Moore – NSF Graduate Research Fellowship (Aug 2017)
- Other metrics
  - Publications: 154
  - Presentations: 104
- Video Presentation at SC18
- Center Webpage
  - http://www.eng.ufl.edu/ccmt/

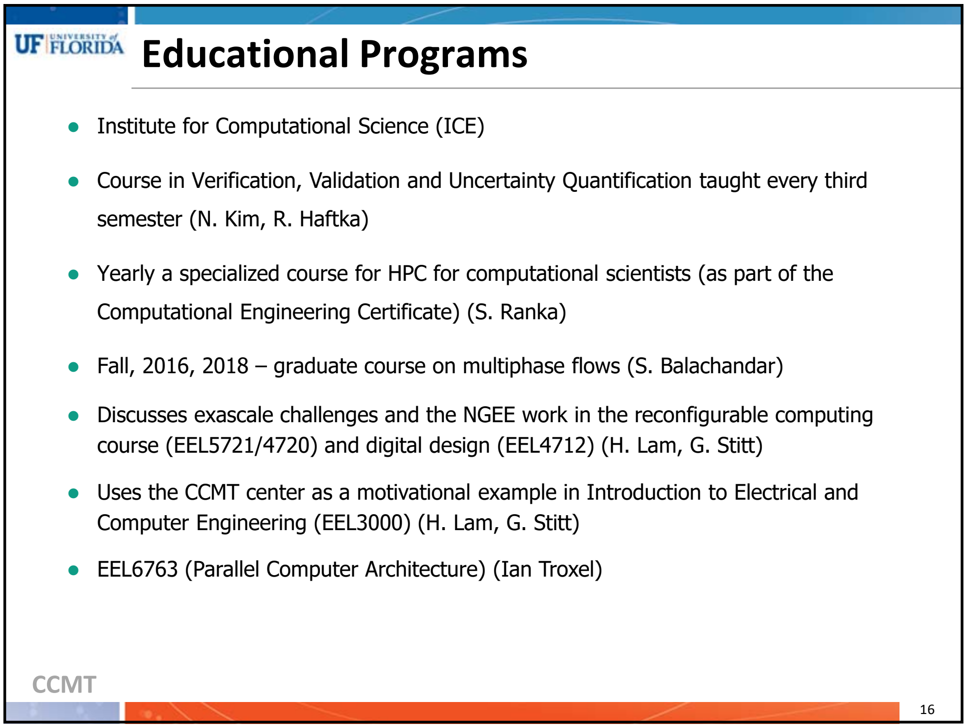

Educational Programs

- Institute for Computational Science (ICE)
- Course in Verification, Validation and Uncertainty Quantification taught every third semester (N. Kim, R. Haftka)
- Yearly a specialized course for HPC for computational scientists (as part of the Computational Engineering Certificate) (S. Ranka)
- Fall, 2016, 2018 – graduate course on multiphase flows (S. Balachandar)
- Discusses exascale challenges and the NGEE work in the reconfigurable computing course (EEL5721/4720) and digital design (EEL4712) (H. Lam, G. Stitt)
- Uses the CCMT center as a motivational example in Introduction to Electrical and Computer Engineering (EEL3000) (H. Lam, G. Stitt)
- EEL6763 (Parallel Computer Architecture) (Ian Troxel)

UF UNIVERSITY of FLORIDA

Management: Tasks and Teams

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

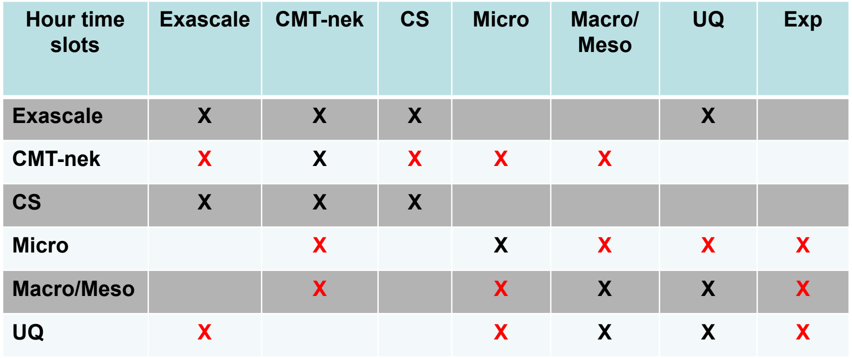

Hour time slots	Exascale	CMT-nek	CS	Micro	Macro/ Meso	UQ	Exp
Exascale	X	X	X			X	
CMT-nek	X	X	X	X	X		
CS	X	X	X				
Micro		X		X	X	X	X
Macro/Meso		X		X	X	X	X
UQ	X			X	X	X	X

- Weekly interactions (black); Regular interactions (red)
- Teams include students, staff, and faculty
- All staff and large number of graduate students located on 2nd floor of PS&T Building
- Construction in PS&T Building to add 6 new office spaces for students
- All meetings held in PS&T Building

UF UNIVERSITY of FLORIDA

4th Annual CCMT Picnic (March, 2018)

Outline

- Background
- Y5 accomplishments
- Y6 plans - Integration and timeline

Prediction Metrics

PM-1: Blast Wave Location

PM-2: Particle Front Location

PM-3: Number of Instability Waves

PM-4: Amplitude of Instability Waves

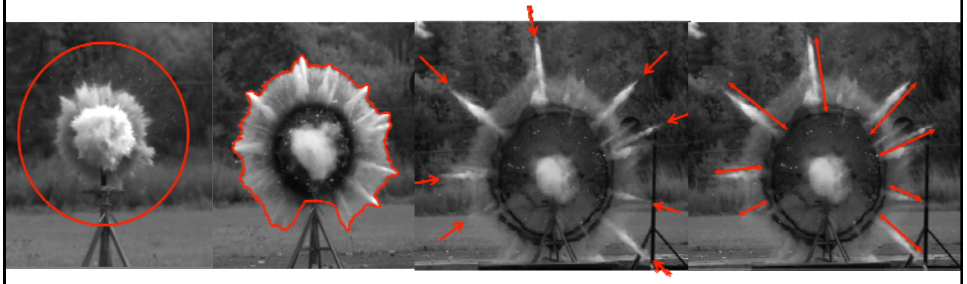

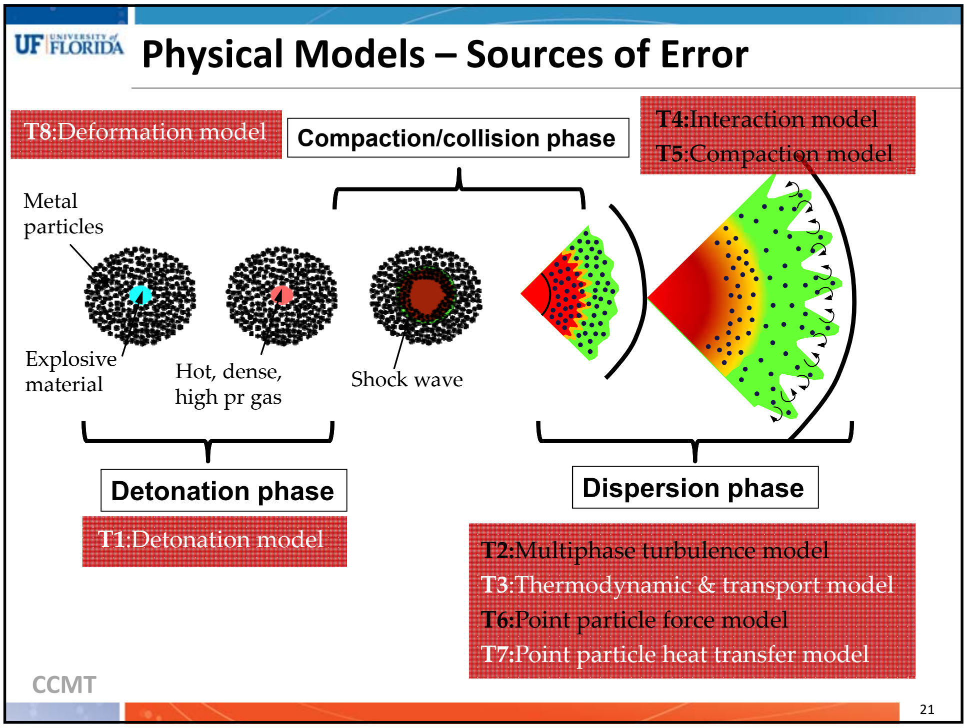

UF UNIVERSITY of FLORIDA

Physical Models – Sources of Error

T8:Deformation model

Compaction/collision phase

T4: Interaction model
T5:Compaction model

Metal particles

Explosive material

Hot, dense, high pr gas

Shock wave

Detonation phase

T1:Detonation model

Dispersion phase

T2: Multiphase turbulence model
T3:Thermodynamic & transport model
T6: Point particle force model
T7: Point particle heat transfer model

CCMT

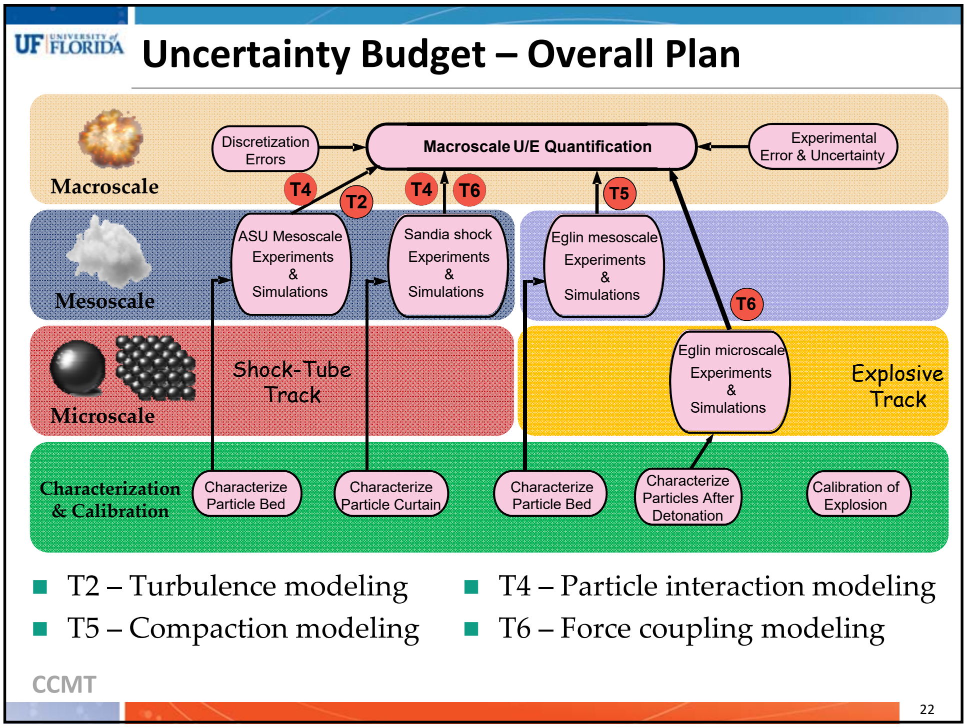

UF UNIVERSITY of FLORIDA

Uncertainty Budget – Overall Plan

Macroscale

Discretization Errors

Macroscale U/E Quantification

Experimental Error & Uncertainty

T4 T2 T4 T6 T5

Mesoscale

ASU Mesoscale Experiments & Simulations

Sandia shock Experiments & Simulations

Eglin mesoscale Experiments & Simulations

T6

Microscale

Shock-Tube Track

Eglin microscale Experiments & Simulations

Explosive Track

Characterization & Calibration

Characterize Particle Bed

Characterize Particle Curtain

Characterize Particle Bed

Characterize Particles After Detonation

Calibration of Explosion

- T2 – Turbulence modeling
- T5 – Compaction modeling
- T4 – Particle interaction modeling
- T6 – Force coupling modeling

CCMT

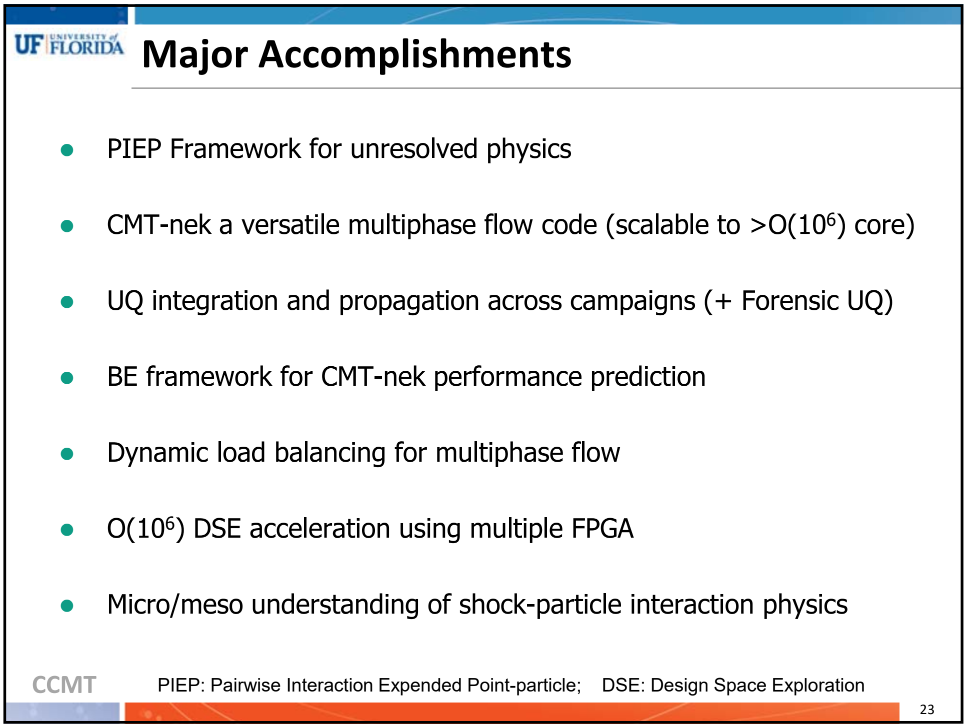

Major Accomplishments

- PIEP Framework for unresolved physics
- CMT-nek a versatile multiphase flow code (scalable to >O(10^6) core)
- UQ integration and propagation across campaigns (+ Forensic UQ)
- BE framework for CMT-nek performance prediction
- Dynamic load balancing for multiphase flow
- O(10^6) DSE acceleration using multiple FPGA
- Micro/meso understanding of shock-particle interaction physics

PIEP: Pairwise Interaction Expended Point-particle; DSE: Design Space Exploration

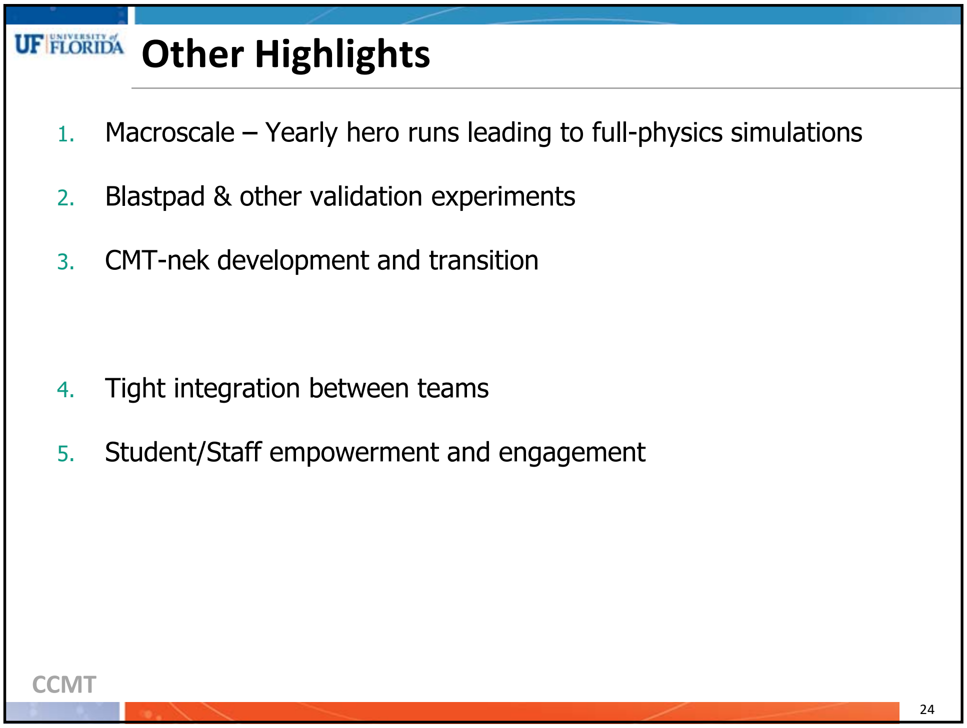

Other Highlights

1. Macroscale – Yearly hero runs leading to full-physics simulations
2. Blastpad & other validation experiments
3. CMT-nek development and transition

4. Tight integration between teams
5. Student/Staff empowerment and engagement

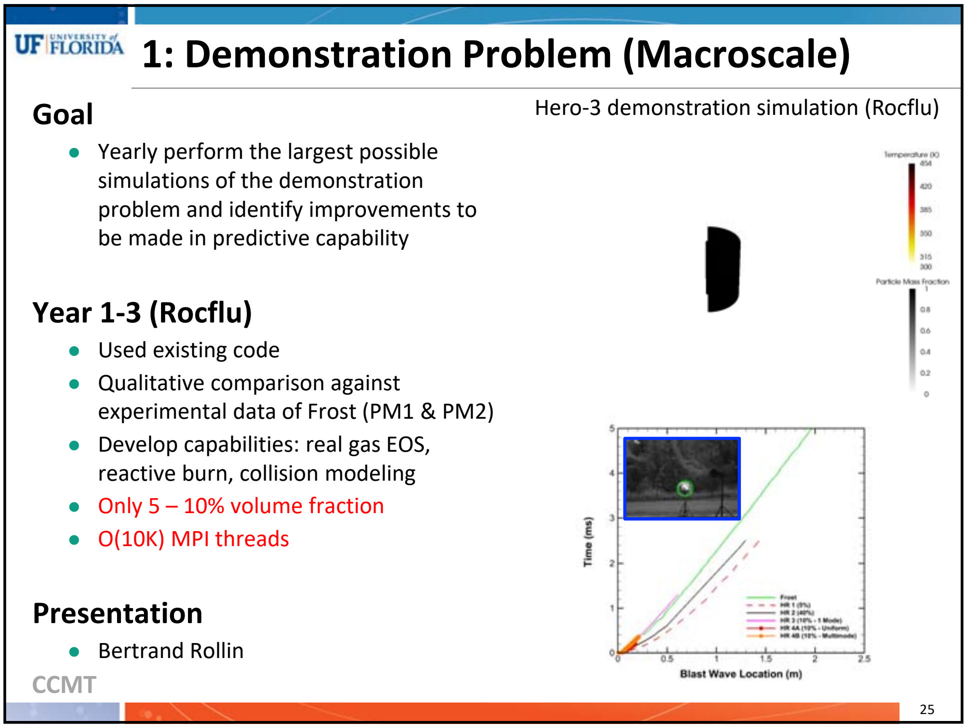

1: Demonstration Problem (Macroscale)

Hero-3 demonstration simulation (Rocflu)

Goal

- Yearly perform the largest possible simulations of the demonstration problem and identify improvements to be made in predictive capability

Year 1-3 (Rocflu)

- Used existing code
- Qualitative comparison against experimental data of Frost (PM1 & PM2)
- Develop capabilities: real gas EOS, reactive burn, collision modeling
- Only 5 – 10% volume fraction
- O(10K) MPI threads

Presentation

- Bertrand Rollin

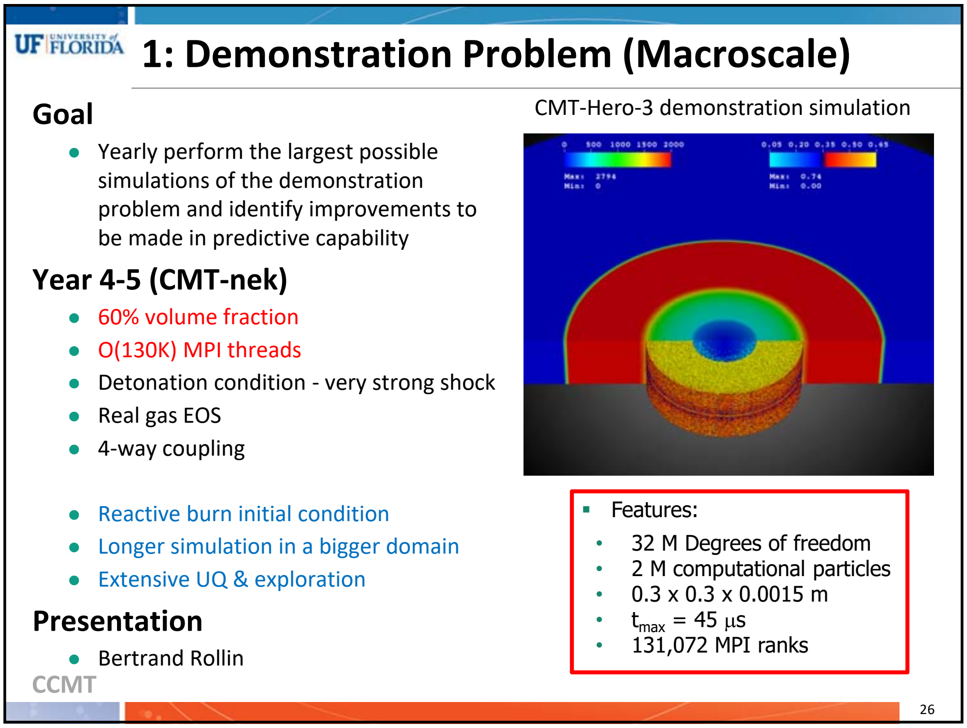

1: Demonstration Problem (Macroscale)

CMT-Hero-3 demonstration simulation

Goal

- Yearly perform the largest possible simulations of the demonstration problem and identify improvements to be made in predictive capability

Year 4-5 (CMT-nek)

- 60% volume fraction
- O(130K) MPI threads
- Detonation condition - very strong shock
- Real gas EOS
- 4-way coupling

- Reactive burn initial condition
- Longer simulation in a bigger domain
- Extensive UQ & exploration

- Features:
  - 32 M Degrees of freedom
  - 2 M computational particles
  - 0.3 x 0.3 x 0.0015 m
  - t_{max} = 45 μs
  - 131,072 MPI ranks

Presentation

- Bertrand Rollin

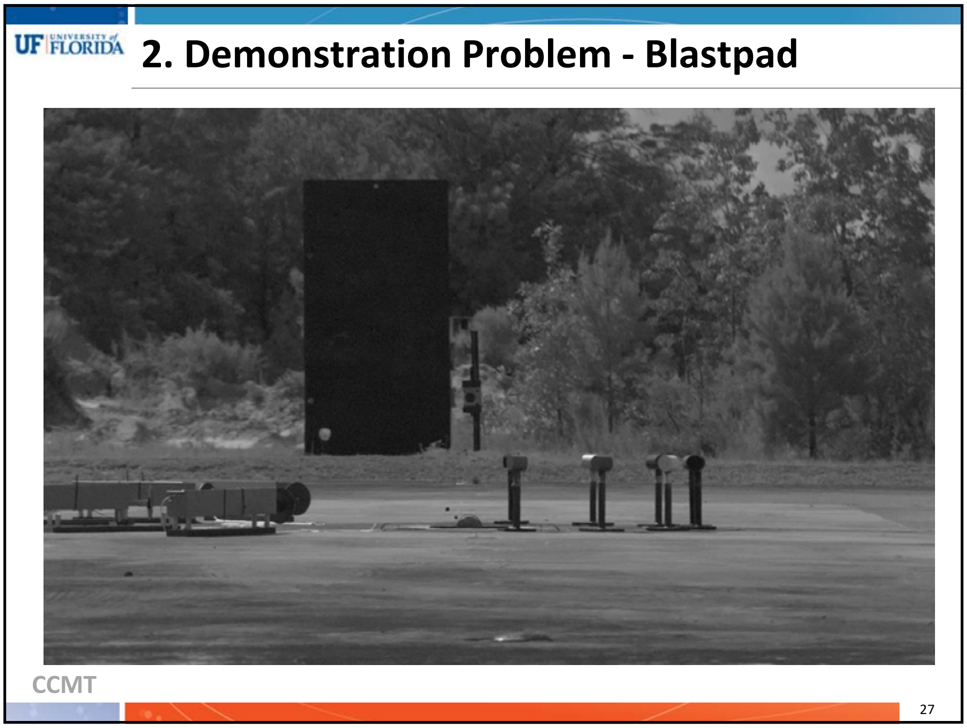

2. Demonstration Problem - Blastpad

CCMT

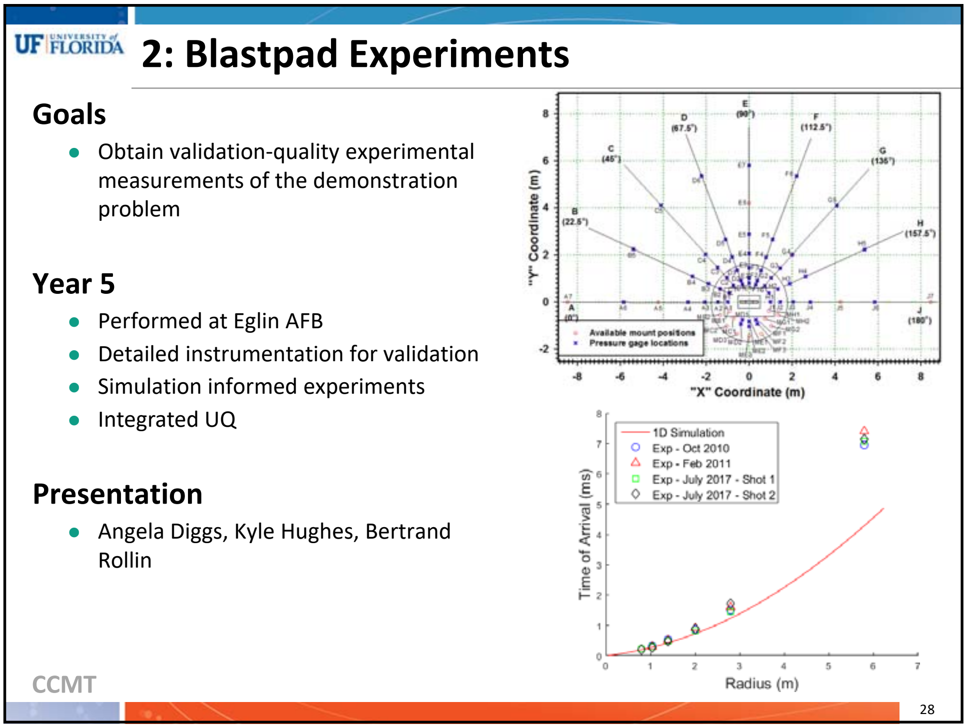

2: Blastpad Experiments

Goals

- Obtain validation-quality experimental measurements of the demonstration problem

Year 5

- Performed at Eglin AFB
- Detailed instrumentation for validation
- Simulation informed experiments
- Integrated UQ

Presentation

- Angela Diggs, Kyle Hughes, Bertrand Rollin

CCMT

2: Blastpad Experiments

Input Parameter	Source
Explosive length [mm]	AFRL measurement
Explosive diameter [mm]	AFRL measurement at 5 locations
Explosive density [kg/m^3]	AFRL calculation
Explosive quality	AFRL X-ray
Particle diameter [mm]	CCMT measurement
Particle density [kg/m^3]	CCMT measurement
Particle volume fraction	AFRL calculation
Ambient pressure [kPa]	AFRL weather station
Ambient temperature [C]	AFRL weather station
Probe locations [m]	CCMT measurement

SEM of single steel particle at 1000x zoom.

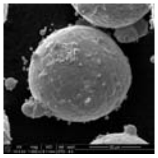

SEM of several steel particles at 100x zoom.

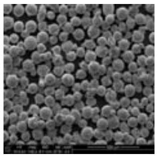

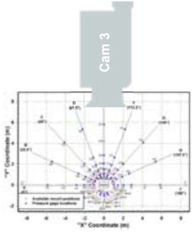

3: CMT-nek Development

Goals

- Co-design an exascale code (CMT-nek) for compressible multiphase turbulence
- Perform micro, meso and demonstration-scale simulations

Year 5

- Microscale simulations with CMT-nek
- Mesoscale simulations with CMT-nek
- Demo simulations with CMT-nek
- CMT-nek in nek5000 repository

Presentation

- Jason Hackl and David Zwick

Mach 3, γ = 1.4,
Color Pressure, contour temperature

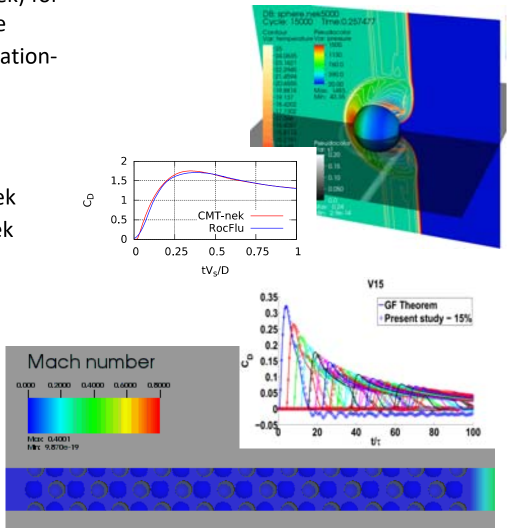

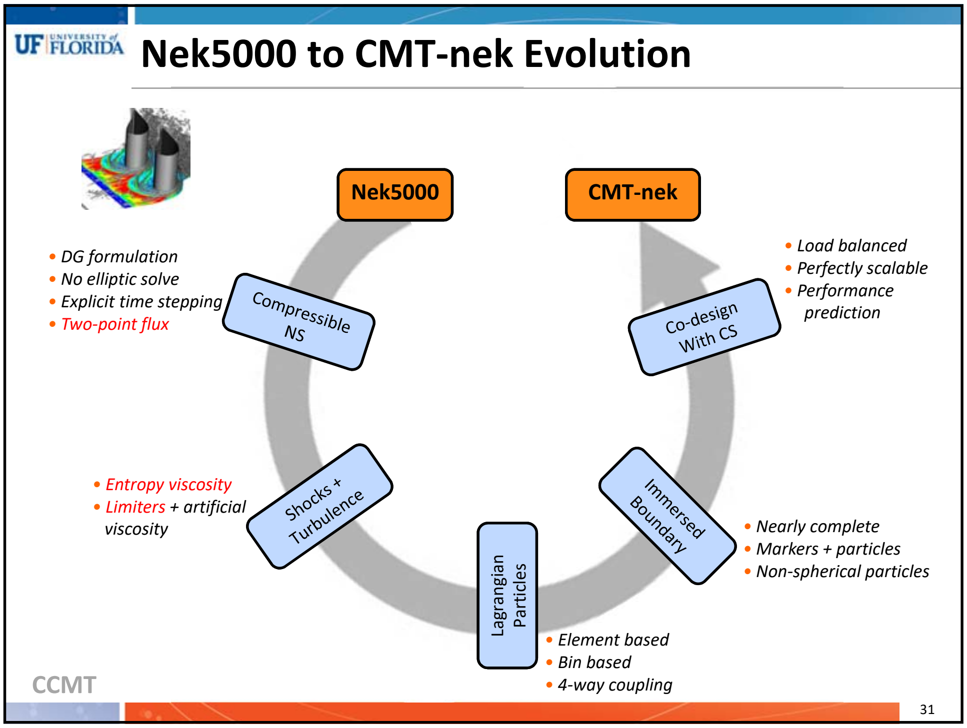

UF UNIVERSITY of FLORIDA

Nek5000 to CMT-nek Evolution

- *DG formulation*
- *No elliptic solve*
- *Explicit time stepping*
- *Two-point flux*

- *Entropy viscosity*
- *Limiters + artificial viscosity*

- *Element based*
- *Bin based*
- *4-way coupling*

- *Nearly complete*
- *Markers + particles*
- *Non-spherical particles*

- *Load balanced*
- *Perfectly scalable*
- *Performance prediction*

CCMT

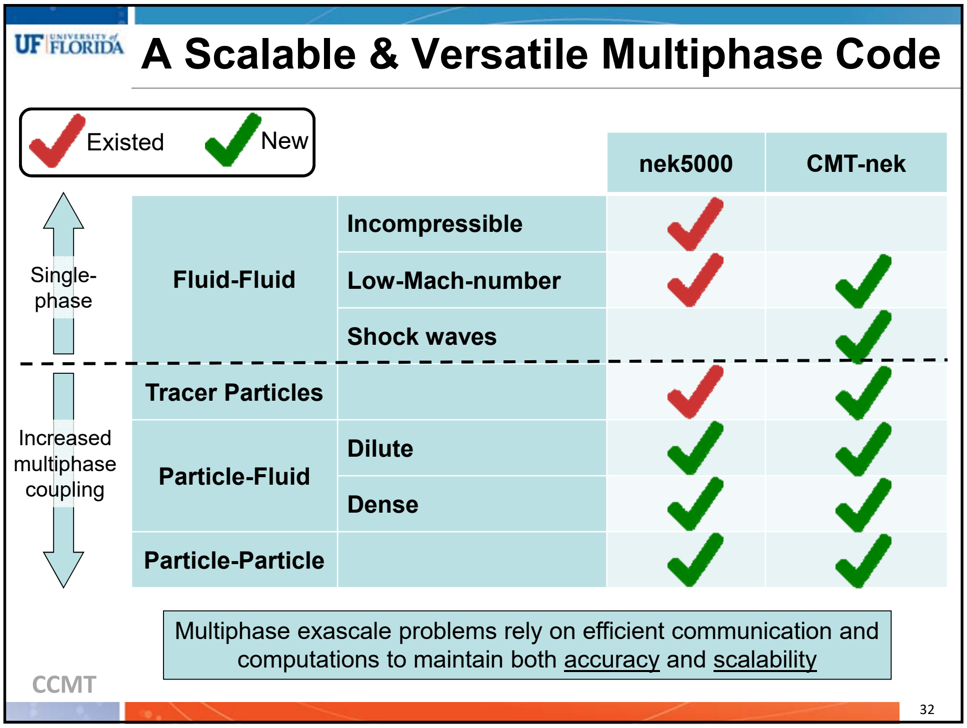

UF UNIVERSITY of FLORIDA

A Scalable & Versatile Multiphase Code

✔ Existed ✅ New

Single-phase

Increased multiphase coupling

		nek5000	CMT-nek
Fluid-Fluid	**Incompressible**	✔	
	Low-Mach-number	✔	✅
	Shock waves		✅
Tracer Particles		✔	✅
Particle-Fluid	**Dilute**	✅	✅
	Dense	✅	✅
Particle-Particle		✅	✅

Multiphase exascale problems rely on efficient communication and computations to maintain both accuracy and scalability

CCMT

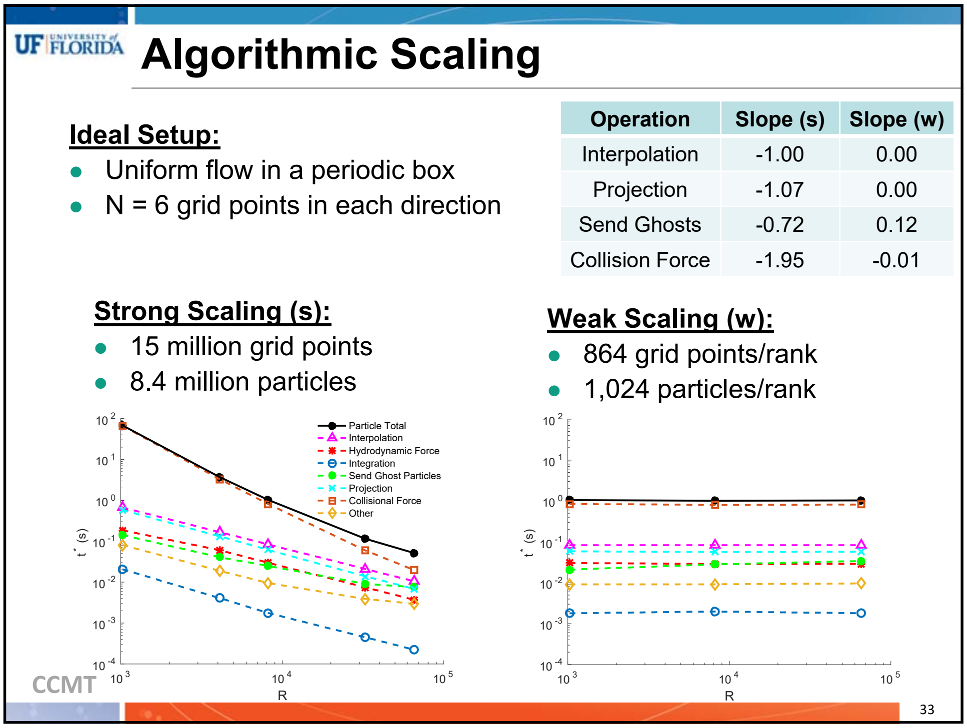

UF UNIVERSITY of FLORIDA

Algorithmic Scaling

Operation	Slope (s)	Slope (w)
Interpolation	-1.00	0.00
Projection	-1.07	0.00
Send Ghosts	-0.72	0.12
Collision Force	-1.95	-0.01

Ideal Setup:

- Uniform flow in a periodic box
- N = 6 grid points in each direction

Strong Scaling (s):

- 15 million grid points
- 8.4 million particles

Weak Scaling (w):

- 864 grid points/rank
- 1,024 particles/rank

CCMT

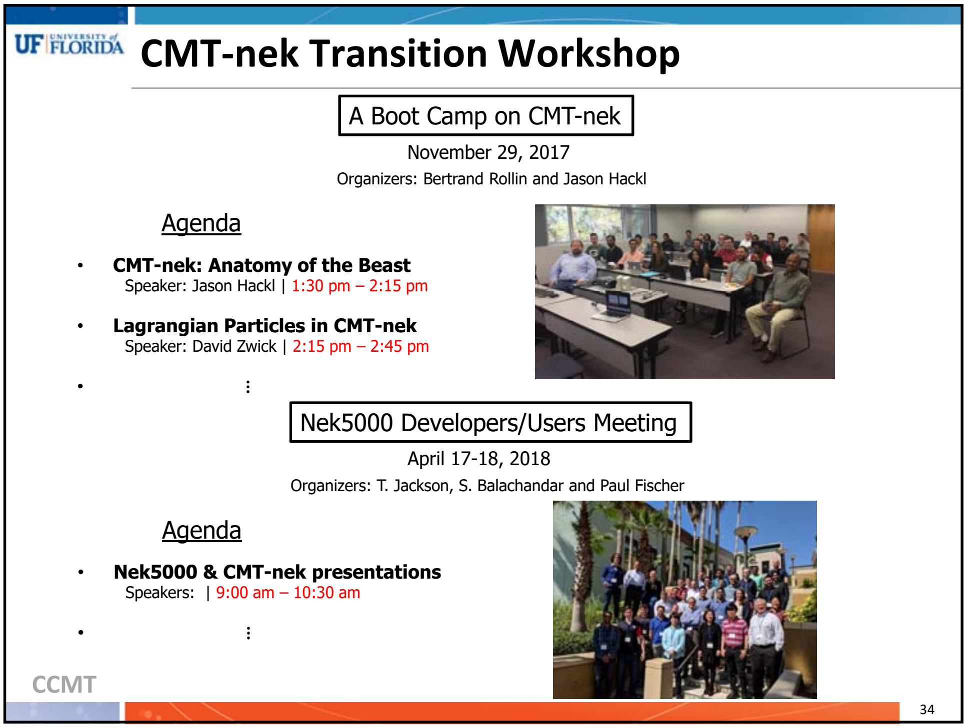

UF UNIVERSITY of FLORIDA

CMT-nek Transition Workshop

A Boot Camp on CMT-nek

November 29, 2017

Organizers: Bertrand Rollin and Jason Hackl

Agenda

- **CMT-nek: Anatomy of the Beast**
  Speaker: Jason Hackl | 1:30 pm – 2:15 pm
- **Lagrangian Particles in CMT-nek**
  Speaker: David Zwick | 2:15 pm – 2:45 pm
- ⋮

Nek5000 Developers/Users Meeting

April 17-18, 2018

Organizers: T. Jackson, S. Balachandar and Paul Fischer

Agenda

- **Nek5000 & CMT-nek presentations**
  Speakers: | 9:00 am – 10:30 am
- ⋮

CCMT

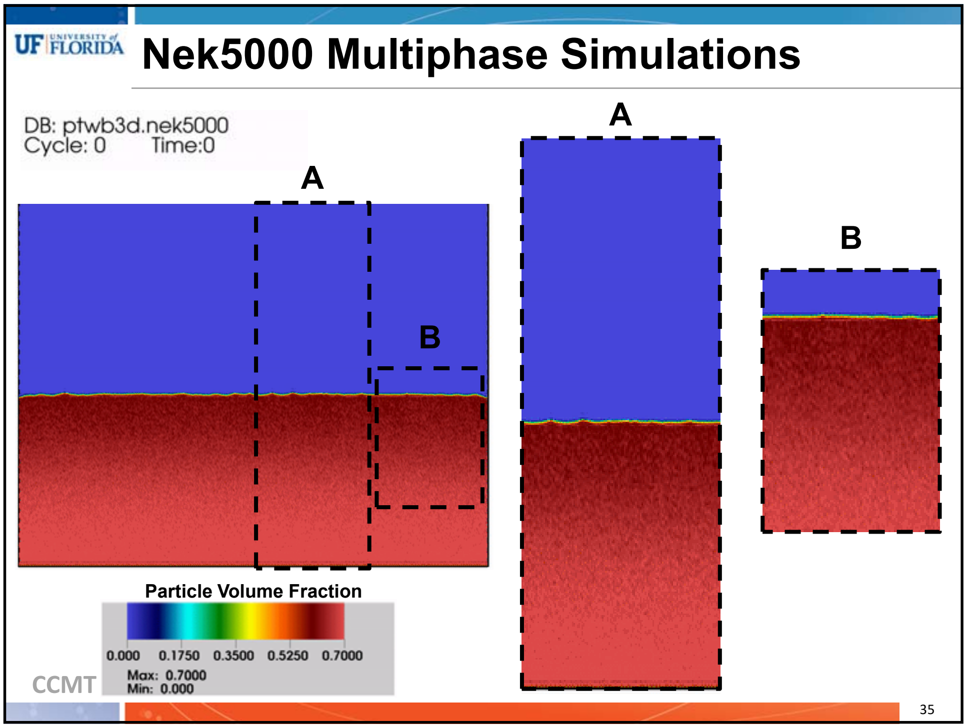

Nek5000 Multiphase Simulations

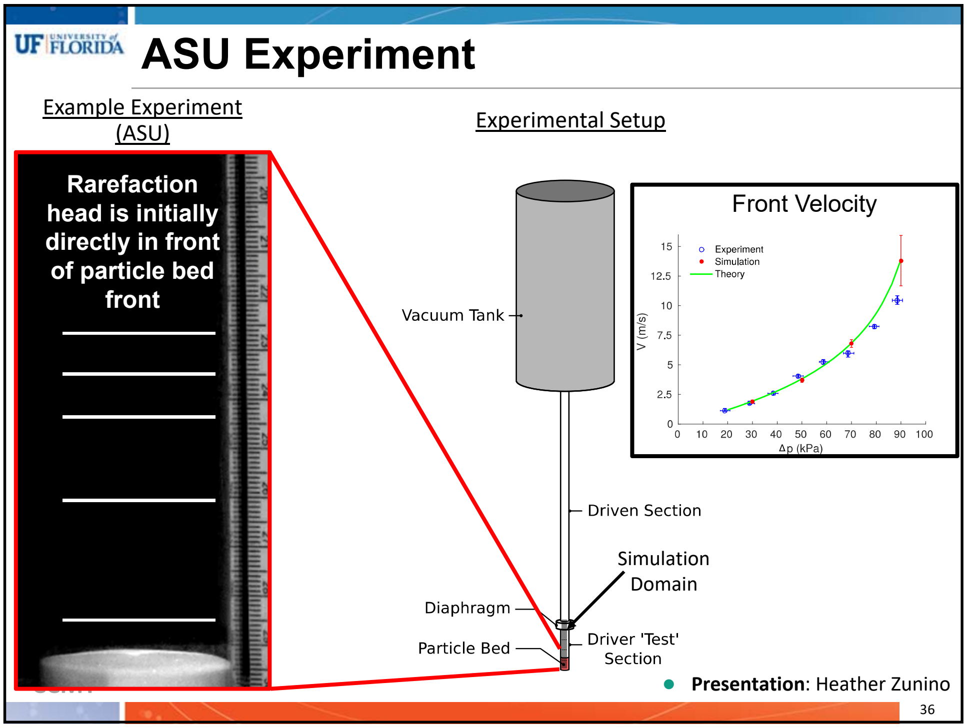

ASU Experiment

Example Experiment (ASU)

Experimental Setup

Rarefaction head is initially directly in front of particle bed front

- **Presentation**: Heather Zunino

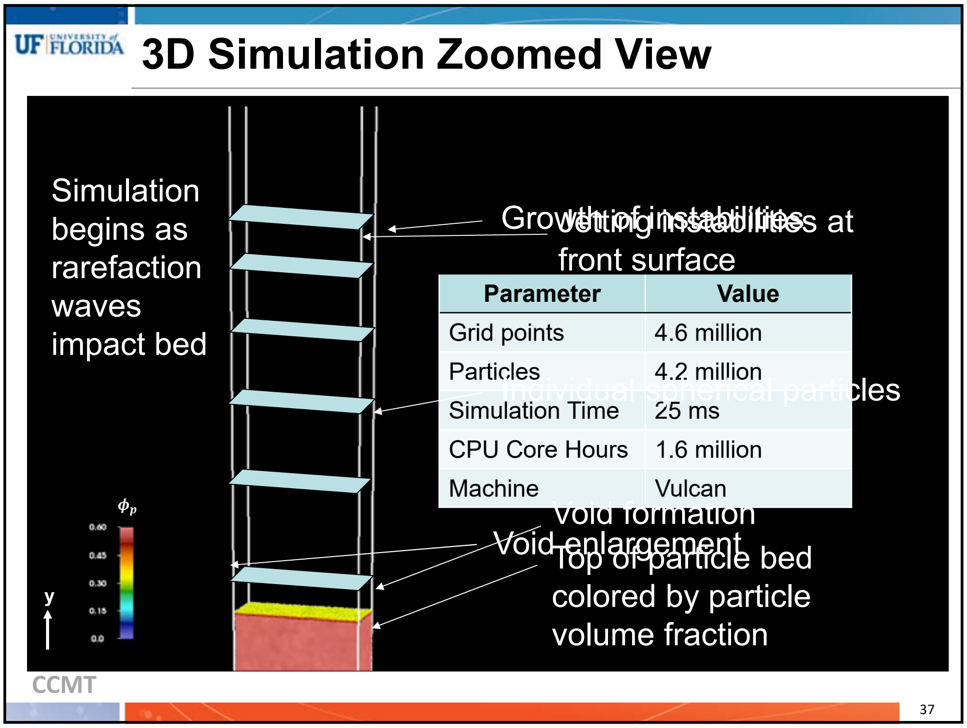

3D Simulation Zoomed View

Simulation begins as rarefaction waves impact bed

Growth of instabilities

Jetting instabilities at front surface

Parameter	Value
Grid points	4.6 million
Particles	4.2 million
Simulation Time	25 ms
CPU Core Hours	1.6 million
Machine	Vulcan

Individual spherical particles

Void formation

Void enlargement

Top of particle bed colored by particle volume fraction

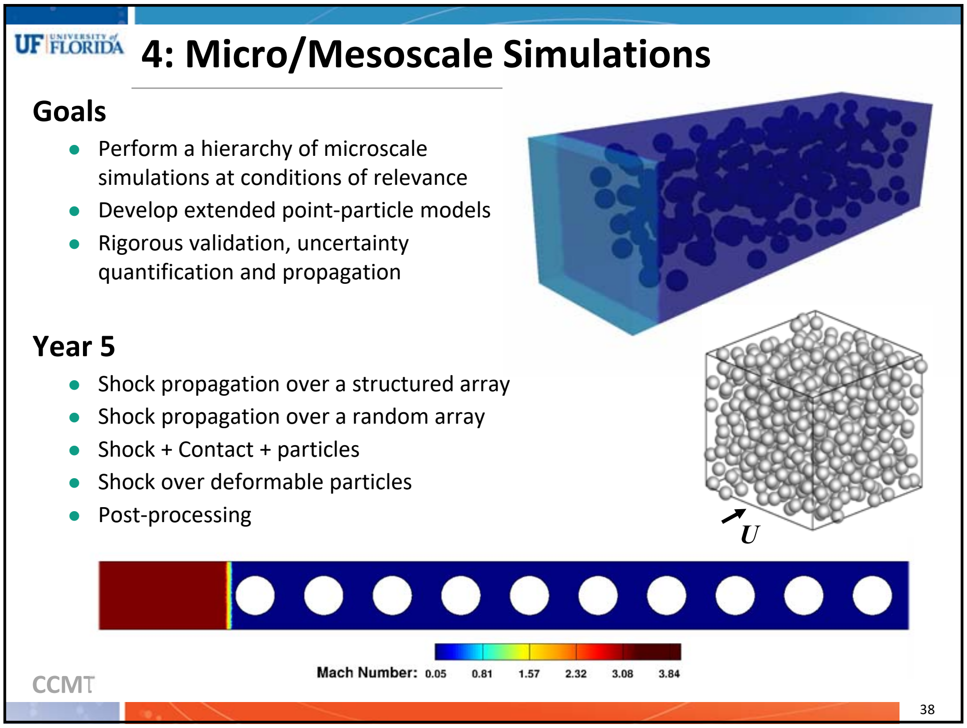

4: Micro/Mesoscale Simulations

Goals

- Perform a hierarchy of microscale simulations at conditions of relevance
- Develop extended point-particle models
- Rigorous validation, uncertainty quantification and propagation

Year 5

- Shock propagation over a structured array
- Shock propagation over a random array
- Shock + Contact + particles
- Shock over deformable particles
- Post-processing

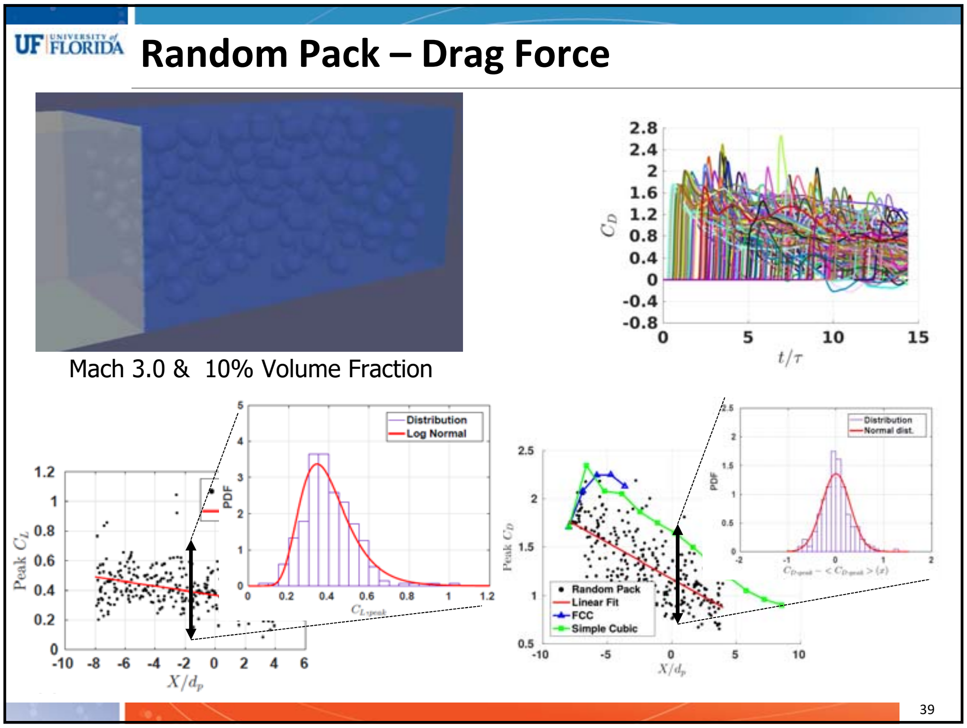

Random Pack – Drag Force

Mach 3.0 & 10% Volume Fraction

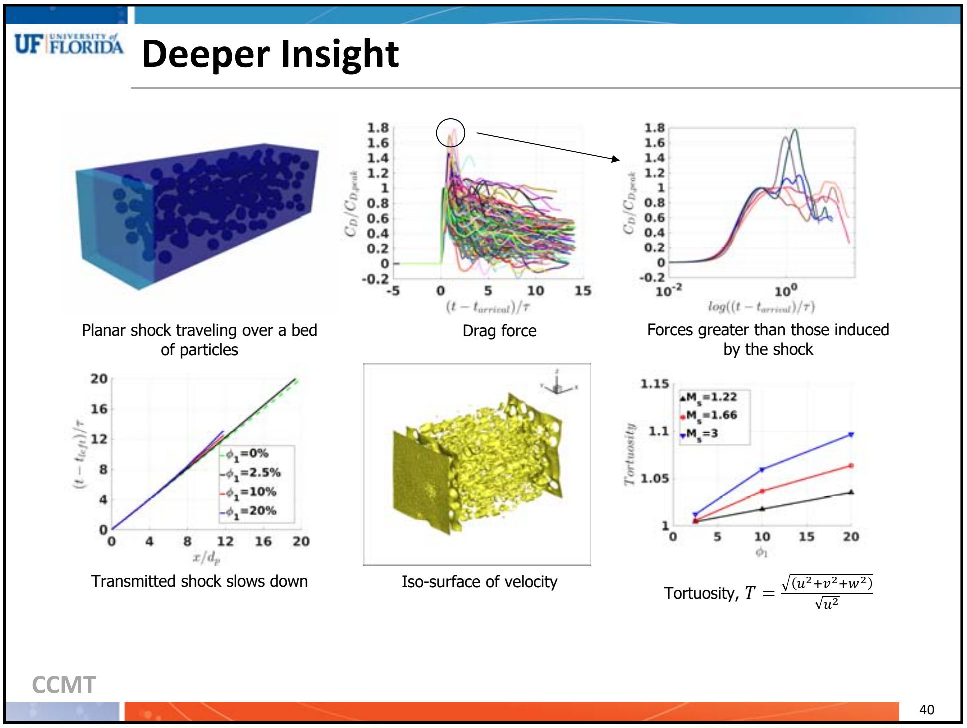

Deeper Insight

Planar shock traveling over a bed of particles

Drag force

Forces greater than those induced by the shock

Transmitted shock slows down

Iso-surface of velocity

Tortuosity, $T = \frac{\sqrt{(u^2+v^2+w^2)}}{\sqrt{u^2}}$

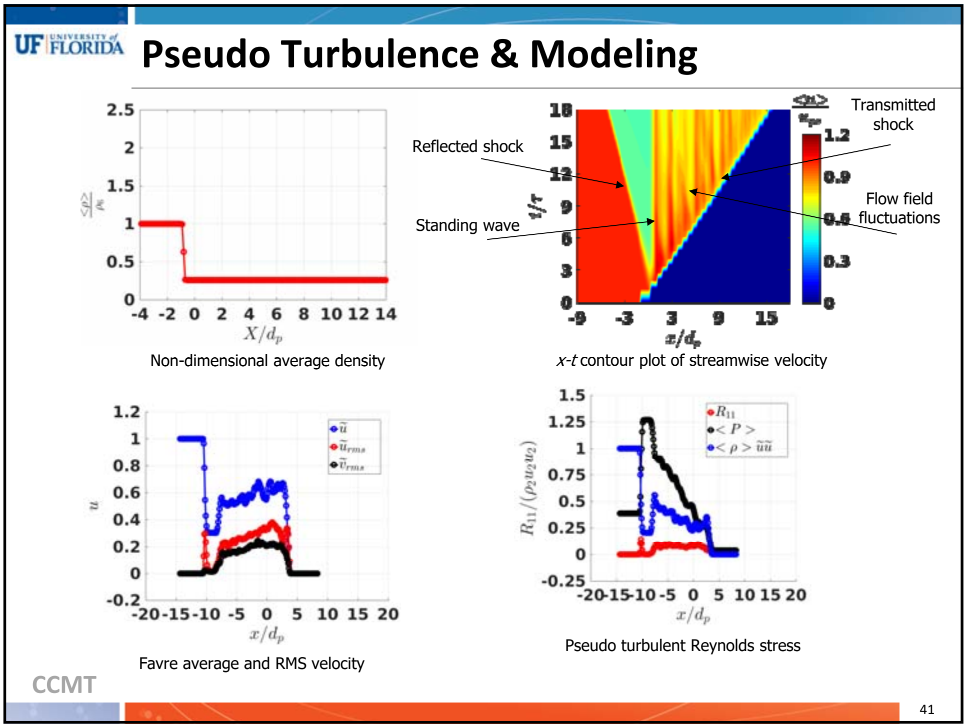

Pseudo Turbulence & Modeling

Reflected shock

Transmitted shock

Standing wave

Flow field fluctuations

Non-dimensional average density

x-t contour plot of streamwise velocity

Favre average and RMS velocity

Pseudo turbulent Reynolds stress

CCMT

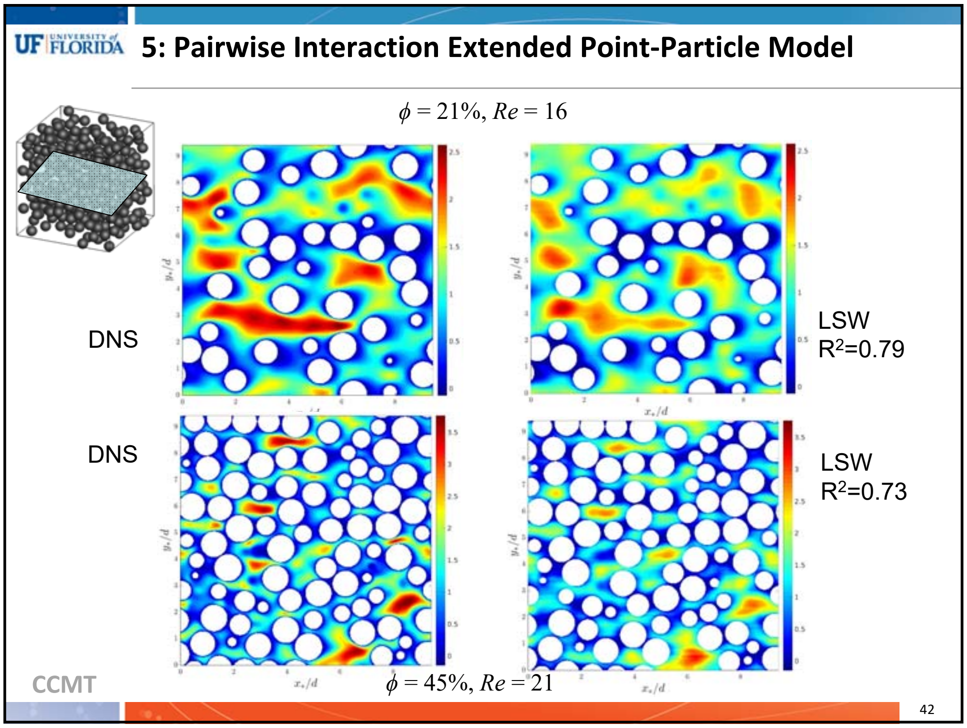

5: Pairwise Interaction Extended Point-Particle Model

$\phi = 21\%, Re = 16$

DNS

LSW
R^2=0.79

DNS

LSW
R^2=0.73

$\phi = 45\%, Re = 21$

CCMT

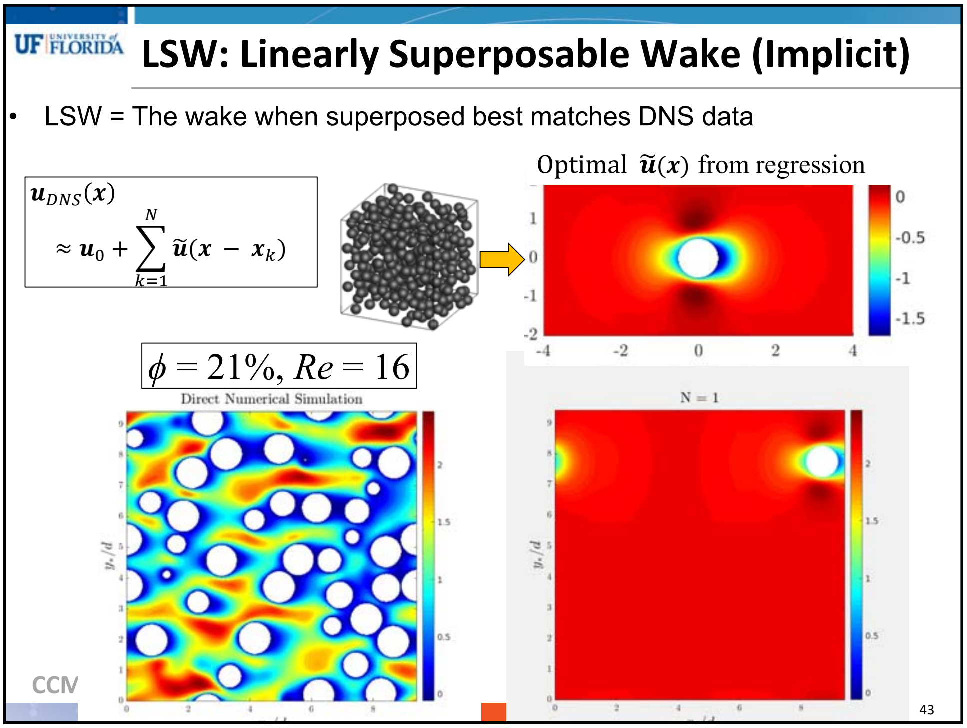

LSW: Linearly Superposable Wake (Implicit)

- LSW = The wake when superposed best matches DNS data

$$\boldsymbol{u}_{DNS}(\boldsymbol{x}) \approx \boldsymbol{u}_0 + \sum_{k=1}^{N} \tilde{\boldsymbol{u}}(\boldsymbol{x} - \boldsymbol{x}_k)$$

Optimal $\tilde{\boldsymbol{u}}(\boldsymbol{x})$ from regression

$\phi = 21\%, Re = 16$

Direct Numerical Simulation

N = 1

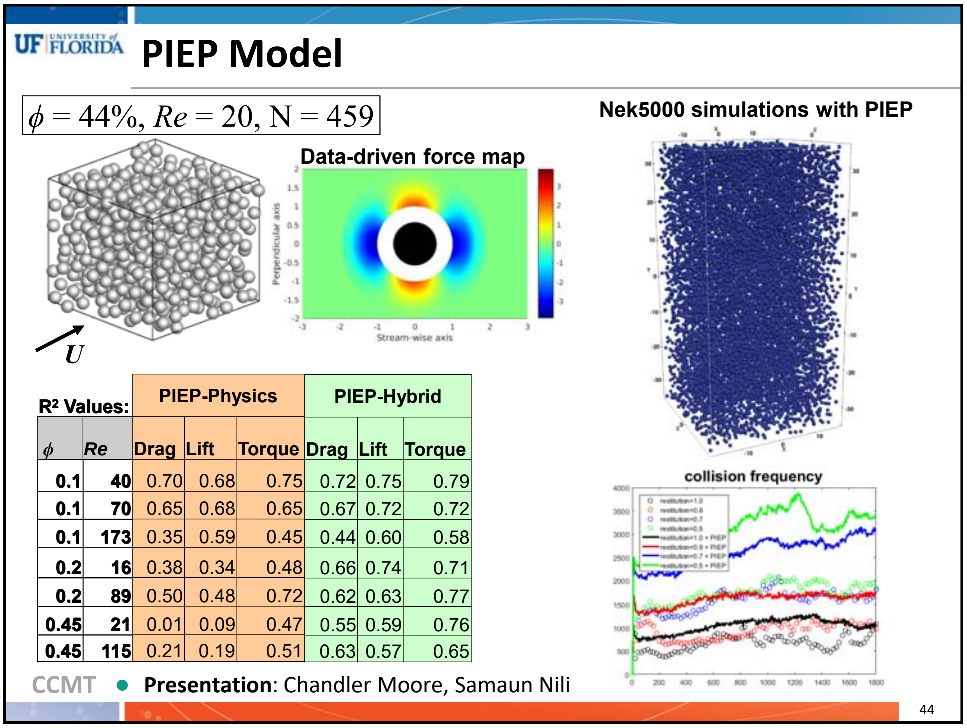

PIEP Model

$\phi = 44\%, Re = 20, \mathrm{N} = 459$

Nek5000 simulations with PIEP

Data-driven force map

R² Values:		**PIEP-Physics**			**PIEP-Hybrid**		
ϕ	*Re*	**Drag**	**Lift**	**Torque**	**Drag**	**Lift**	**Torque**
0.1	**40**	0.70	0.68	0.75	0.72	0.75	0.79
0.1	**70**	0.65	0.68	0.65	0.67	0.72	0.72
0.1	**173**	0.35	0.59	0.45	0.44	0.60	0.58
0.2	**16**	0.38	0.34	0.48	0.66	0.74	0.71
0.2	**89**	0.50	0.48	0.72	0.62	0.63	0.77
0.45	**21**	0.01	0.09	0.47	0.55	0.59	0.76
0.45	**115**	0.21	0.19	0.51	0.63	0.57	0.65

collision frequency

- **Presentation**: Chandler Moore, Samaun Nili

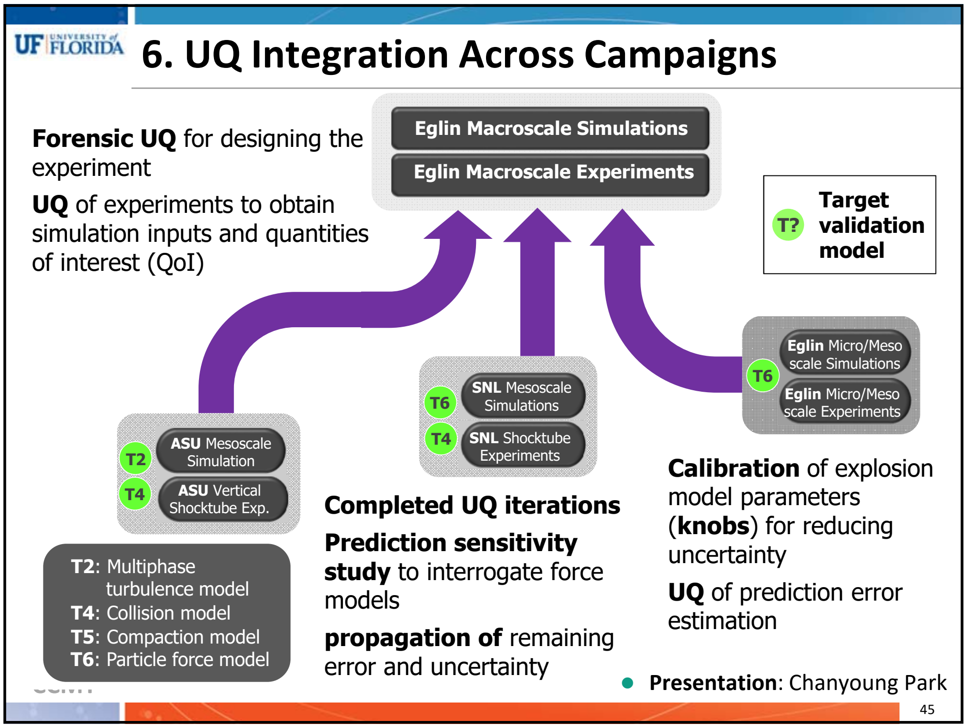

6. UQ Integration Across Campaigns

Forensic UQ for designing the experiment

UQ of experiments to obtain simulation inputs and quantities of interest (QoI)

Eglin Macroscale Simulations

Eglin Macroscale Experiments

T? **Target validation model**

T2 **ASU** Mesoscale Simulation

T4 **ASU** Vertical Shocktube Exp.

T6 **SNL** Mesoscale Simulations

T4 **SNL** Shocktube Experiments

T6 **Eglin** Micro/Meso scale Simulations

Eglin Micro/Meso scale Experiments

T2: Multiphase turbulence model
T4: Collision model
T5: Compaction model
T6: Particle force model

Completed UQ iterations

Prediction sensitivity study to interrogate force models

propagation of remaining error and uncertainty

Calibration of explosion model parameters (**knobs**) for reducing uncertainty

UQ of prediction error estimation

- **Presentation**: Chanyoung Park

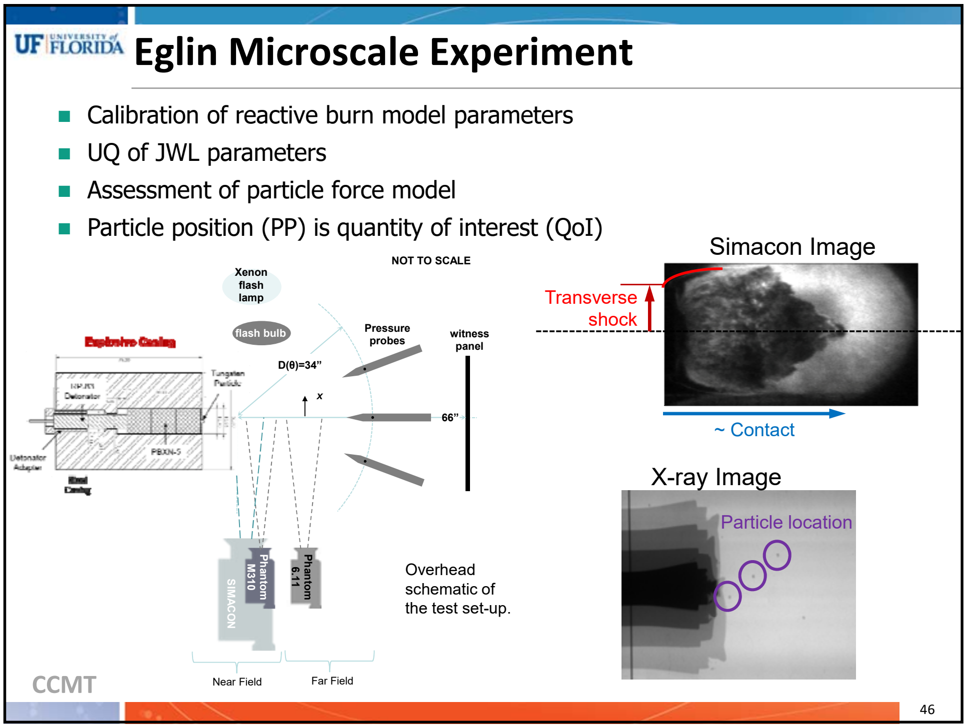

Eglin Microscale Experiment

- Calibration of reactive burn model parameters
- UQ of JWL parameters
- Assessment of particle force model
- Particle position (PP) is quantity of interest (QoI)

Overhead schematic of the test set-up.

CCMT

Calibration - Uncertainty Reduction

Parameter	Quantity	Method
U1: Explosive density	mean = 1795 kg/m^3, std = 2.9 kg/m^3, Bi-modal	Derived
U2: RZ thickness	[0.365, 1.4] mm	Expert opinion
U3: Heat release Q	[11.6, 14.2] MJ/kg	Literature and opinion
U4: Particle diameter	mean = 2.0156 mm, std = 0.0073 mm, Weibull	Direct measurement Micrometer, 52 samples
U5: Particle density	mean = 15540 kg/m^3, std = 250 kg/m^3, Normal	Pycnometer Gas Pycnometer, 12 samples
U6: Initial radial position	[0, 0.254] mm	
U7-11: JWL	A, B, w, R$_1$, R$_2$	Literature and opinion
…	…	…

- Including uncertainties in simulation parameters (knob)
- **Kyle** actively involved in measurements and uncertainty estimation (forensic UQ)

Two Most Probable Detonation Parameter

- Predict **transverse shock position** and **particle position** with the micro simulation
- Transverse shock position and particle position **do not show meaningful difference** for the two most probable parameter estimates
- **Presentation**: Josh Garno

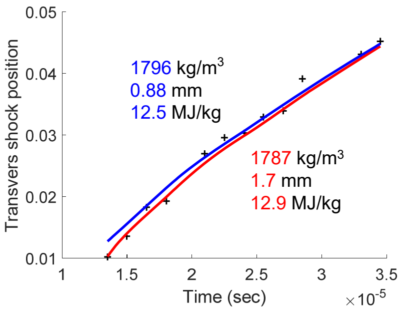

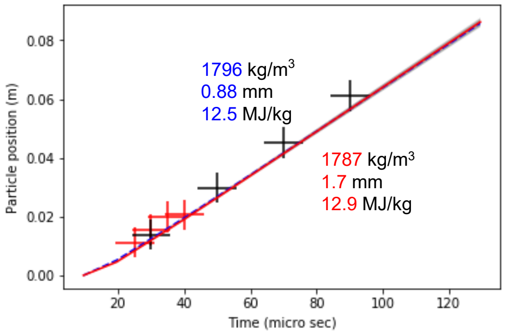

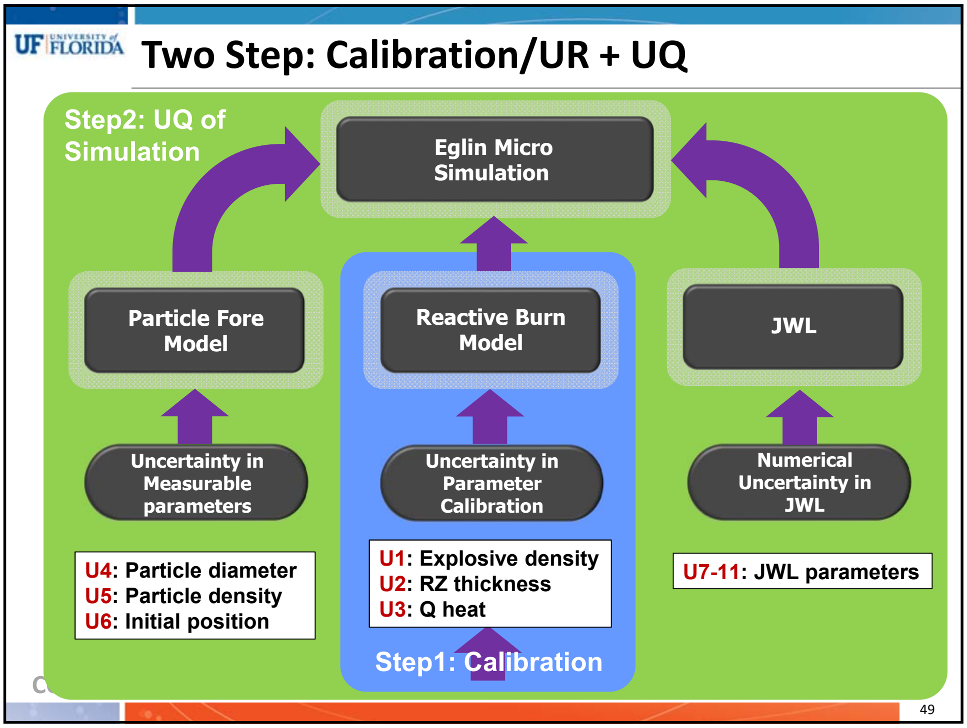

Two Step: Calibration/UR + UQ

Step2: UQ of Simulation

Eglin Micro Simulation

Particle Fore Model

Reactive Burn Model

JWL

Uncertainty in Measurable parameters

Uncertainty in Parameter Calibration

Numerical Uncertainty in JWL

U4: Particle diameter
U5: Particle density
U6: Initial position

U1: Explosive density
U2: RZ thickness
U3: Q heat

U7-11: JWL parameters

Step1: Calibration

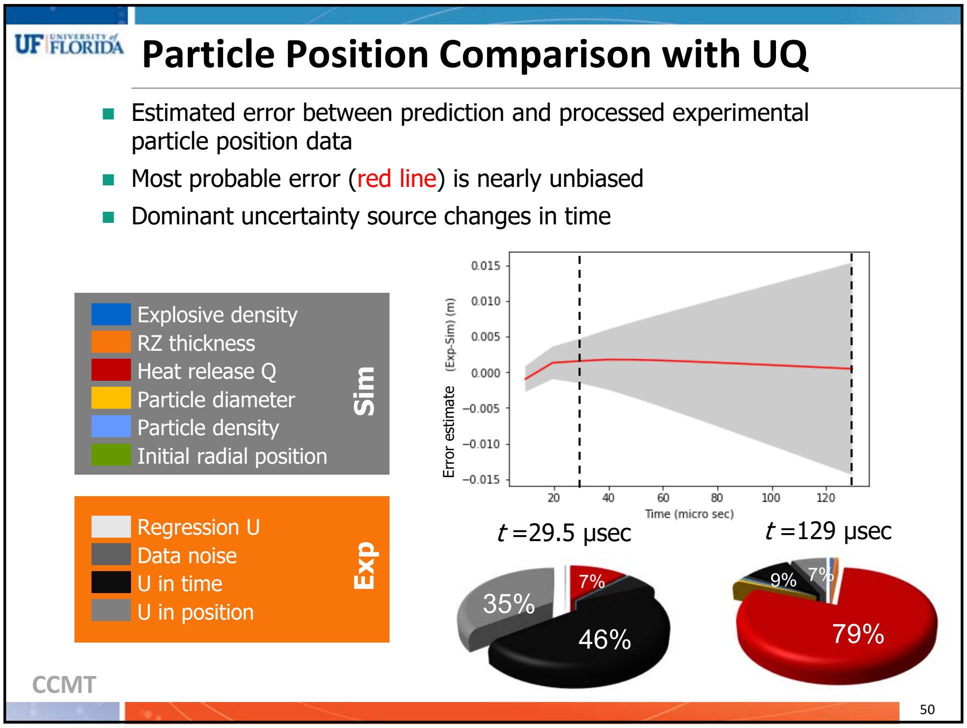

Particle Position Comparison with UQ

- Estimated error between prediction and processed experimental particle position data
- Most probable error (red line) is nearly unbiased
- Dominant uncertainty source changes in time

Sim
🟦 Explosive density
🟧 RZ thickness
🟥 Heat release Q
🟨 Particle diameter
🟦 Particle density
🟩 Initial radial position

Exp
⬜ Regression U
⬛ Data noise
⬛ U in time
⬛ U in position

CCMT

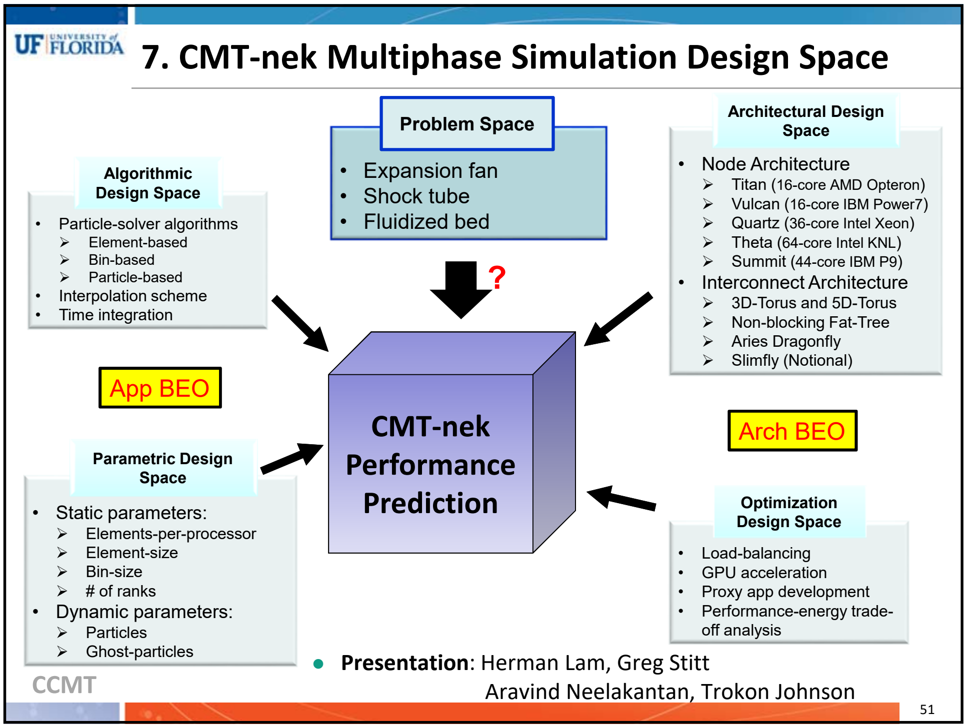

7. CMT-nek Multiphase Simulation Design Space

Problem Space

- Expansion fan
- Shock tube
- Fluidized bed

?

Algorithmic Design Space

- Particle-solver algorithms
  - Element-based
  - Bin-based
  - Particle-based
- Interpolation scheme
- Time integration

App BEO

Parametric Design Space

- Static parameters:
  - Elements-per-processor
  - Element-size
  - Bin-size
  - # of ranks
- Dynamic parameters:
  - Particles
  - Ghost-particles

CMT-nek Performance Prediction

Architectural Design Space

- Node Architecture
  - Titan (16-core AMD Opteron)
  - Vulcan (16-core IBM Power7)
  - Quartz (36-core Intel Xeon)
  - Theta (64-core Intel KNL)
  - Summit (44-core IBM P9)
- Interconnect Architecture
  - 3D-Torus and 5D-Torus
  - Non-blocking Fat-Tree
  - Aries Dragonfly
  - Slimfly (Notional)

Arch BEO

Optimization Design Space

- Load-balancing
- GPU acceleration
- Proxy app development
- Performance-energy trade-off analysis

- **Presentation**: Herman Lam, Greg Stitt
  Aravind Neelakantan, Trokon Johnson

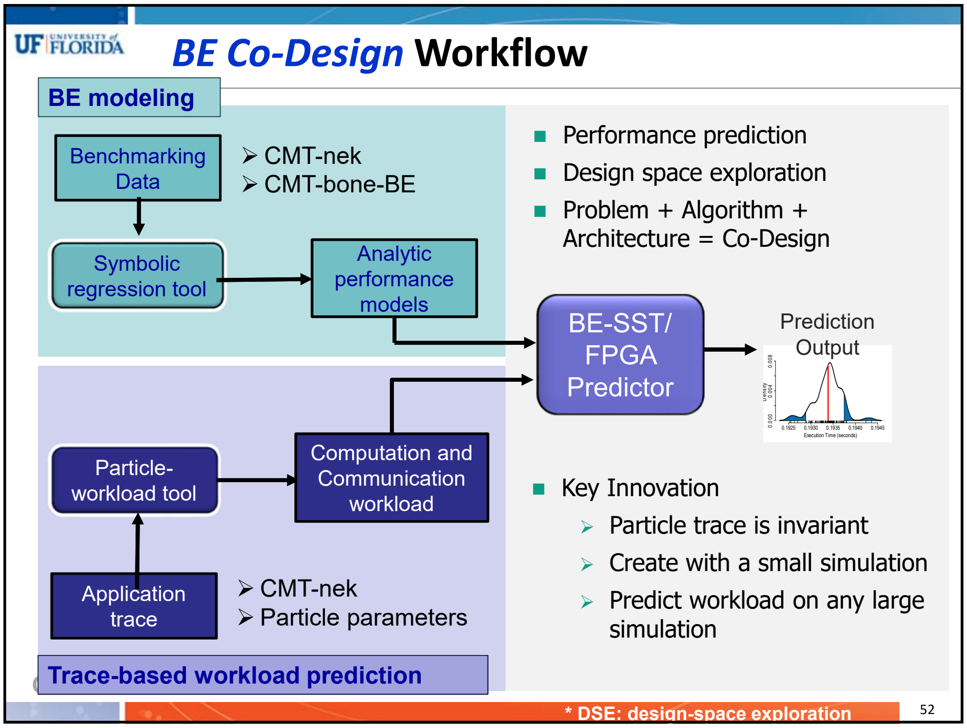

BE Co-Design Workflow

BE modeling

Benchmarking Data

- CMT-nek
- CMT-bone-BE

Symbolic regression tool → Analytic performance models → BE-SST/FPGA Predictor → Prediction Output

Particle-workload tool → Computation and Communication workload → BE-SST/FPGA Predictor

Application trace

- CMT-nek
- Particle parameters

Trace-based workload prediction

- Performance prediction
- Design space exploration
- Problem + Algorithm + Architecture = Co-Design

- Key Innovation
  - Particle trace is invariant
  - Create with a small simulation
  - Predict workload on any large simulation

* DSE: design-space exploration

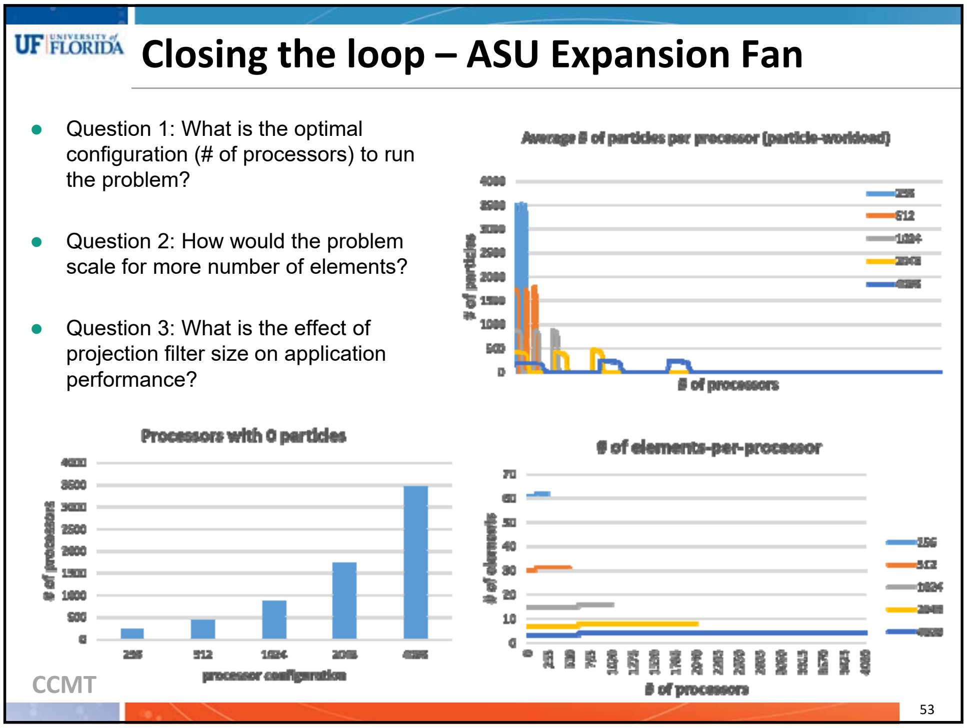

Closing the loop – ASU Expansion Fan

- Question 1: What is the optimal configuration (# of processors) to run the problem?
- Question 2: How would the problem scale for more number of elements?
- Question 3: What is the effect of projection filter size on application performance?

CCMT

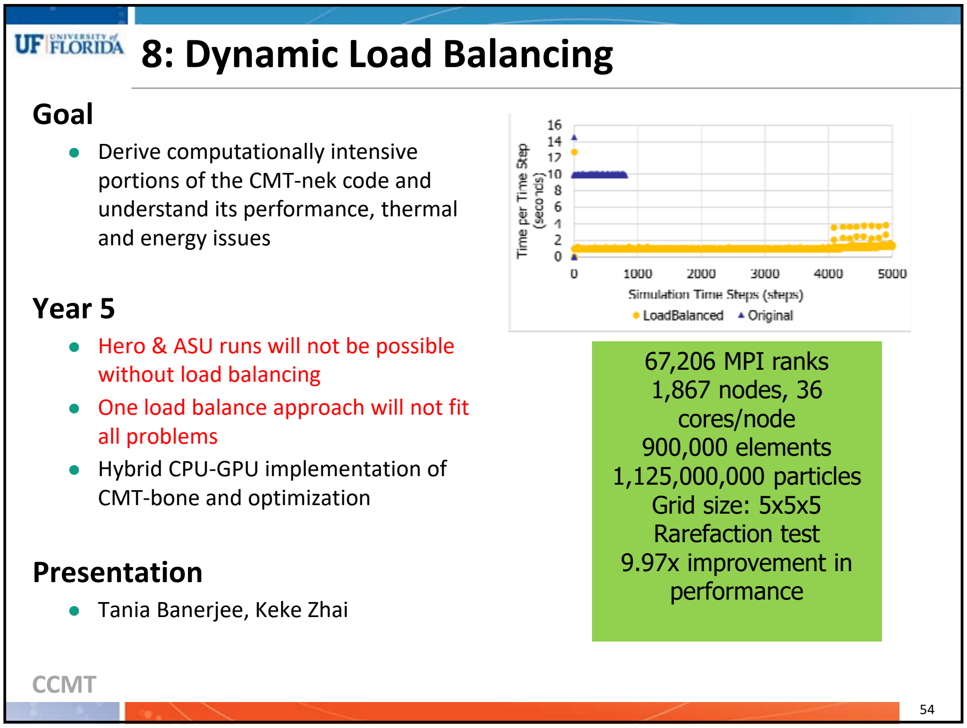

8: Dynamic Load Balancing

Goal

- Derive computationally intensive portions of the CMT-nek code and understand its performance, thermal and energy issues

Year 5

- Hero & ASU runs will not be possible without load balancing
- One load balance approach will not fit all problems
- Hybrid CPU-GPU implementation of CMT-bone and optimization

Presentation

- Tania Banerjee, Keke Zhai

67,206 MPI ranks
1,867 nodes, 36 cores/node
900,000 elements
1,125,000,000 particles
Grid size: 5x5x5
Rarefaction test
9.97x improvement in performance

CCMT

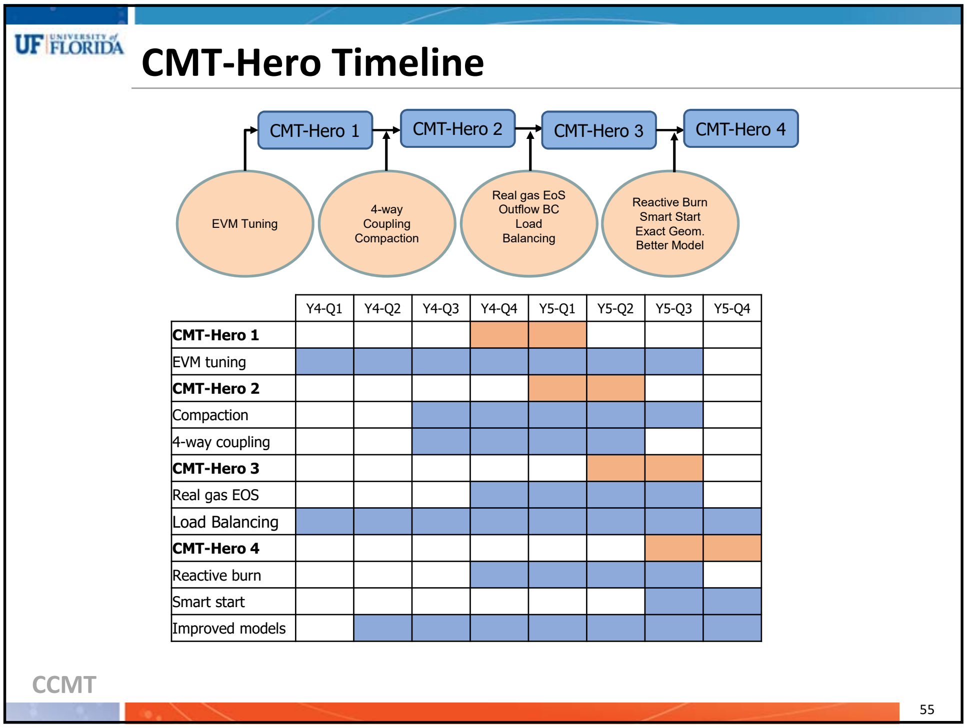

CMT-Hero Timeline

	Y4-Q1	Y4-Q2	Y4-Q3	Y4-Q4	Y5-Q1	Y5-Q2	Y5-Q3	Y5-Q4
CMT-Hero 1				🟧	🟧			
EVM tuning	🟦	🟦	🟦	🟦	🟦	🟦	🟦	
CMT-Hero 2					🟧	🟧		
Compaction			🟦	🟦	🟦	🟦	🟦	
4-way coupling			🟦	🟦	🟦	🟦		
CMT-Hero 3						🟧	🟧	
Real gas EOS				🟦	🟦	🟦	🟦	
Load Balancing	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦
CMT-Hero 4							🟧	🟧
Reactive burn				🟦	🟦	🟦	🟦	
Smart start							🟦	🟦
Improved models		🟦	🟦	🟦	🟦	🟦	🟦	🟦

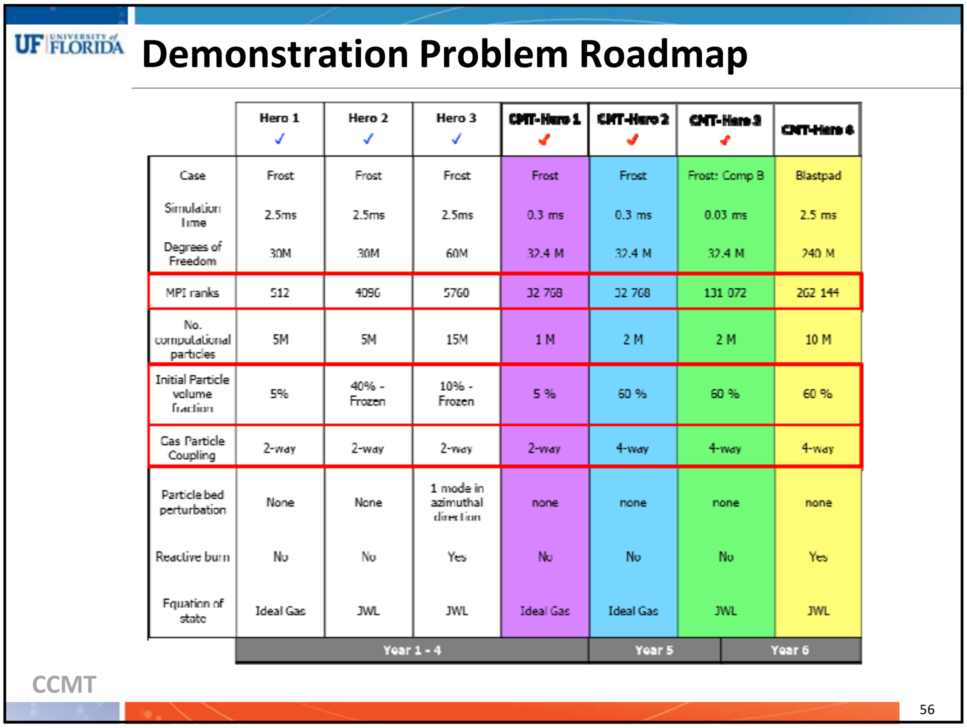

Demonstration Problem Roadmap

	Hero 1 ✓	Hero 2 ✓	Hero 3 ✓	CMT-Hero 1 ✓	CMT-Hero 2 ✓	CMT-Hero 3 ✓	CMT-Hero 4
Case	Frost	Frost	Frost	Frost	Frost	Frost: Comp B	Blastpad
Simulation Time	2.5ms	2.5ms	2.5ms	0.3 ms	0.3 ms	0.03 ms	2.5 ms
Degrees of Freedom	30M	30M	60M	32.4 M	32.4 M	32.4 M	240 M
MPI ranks	512	4096	5760	32 768	32 768	131 072	262 144
No. computational particles	5M	5M	15M	1 M	2 M	2 M	10 M
Initial Particle volume fraction	5%	40% - Frozen	10% - Frozen	5 %	60 %	60 %	60 %
Gas Particle Coupling	2-way	2-way	2-way	2-way	4-way	4-way	4-way
Particle bed perturbation	None	None	1 mode in azimuthal direction	none	none	none	none
Reactive burn	No	No	Yes	No	No	No	Yes
Equation of state	Ideal Gas	JWL	JWL	Ideal Gas	Ideal Gas	JWL	JWL
	Year 1 - 4				Year 5	Year 6	

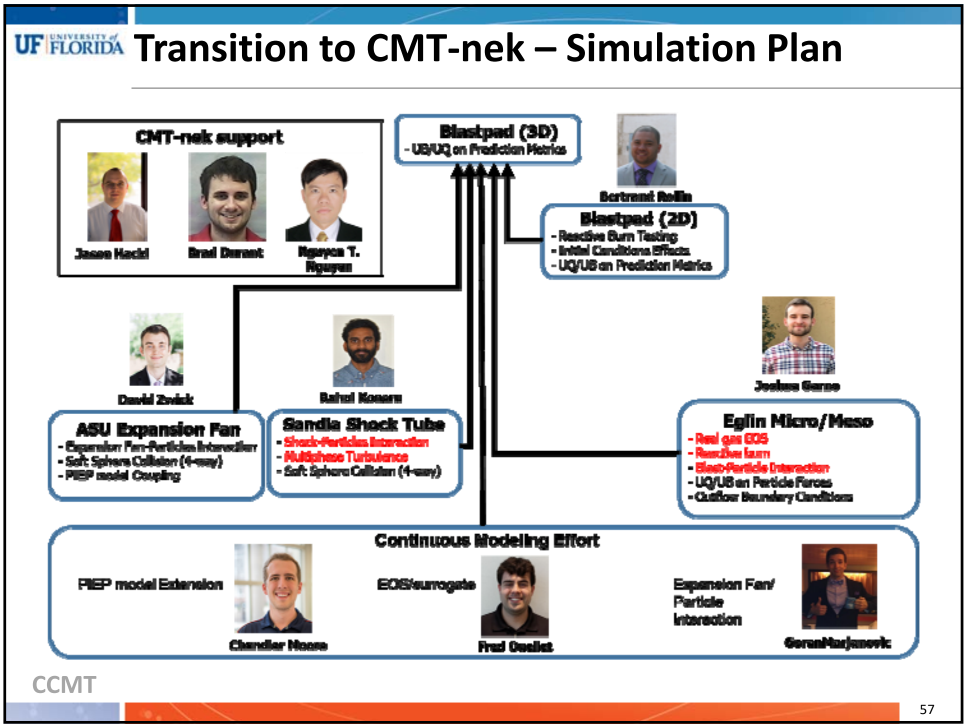

Transition to CMT-nek – Simulation Plan

CMT-nek support

Jason Hackl

Brad Durant

Nguyen T. Nguyen

Blastpad (3D)
- UB/UQ on Prediction Metrics

Bertrand Rollin

Blastpad (2D)
- Reactive Burn Testing
- Initial Conditions Effects
- UQ/UB on Prediction Metrics

David Zwick

ASU Expansion Fan
- Expansion Fan-Particles Interaction
- Soft Sphere Collision (4-way)
- PIEP model Coupling

Rahul Koneru

Sandia Shock Tube
- Shock-Particles Interaction
- Multiphase Turbulence
- Soft Sphere Collision (4-way)

Joshua Garno

Eglin Micro/Meso
- Real gas EOS
- Reactive burn
- Blast-Particle Interaction
- UQ/UB on Particle Forces
- Outflow Boundary Conditions

Continuous Modeling Effort

PIEP model Extension — Chandler Moore

EOS/surrogate — Fred Ouellet

Expansion Fan/ Particle Interaction — GoranMarjanovic

CCMT

57

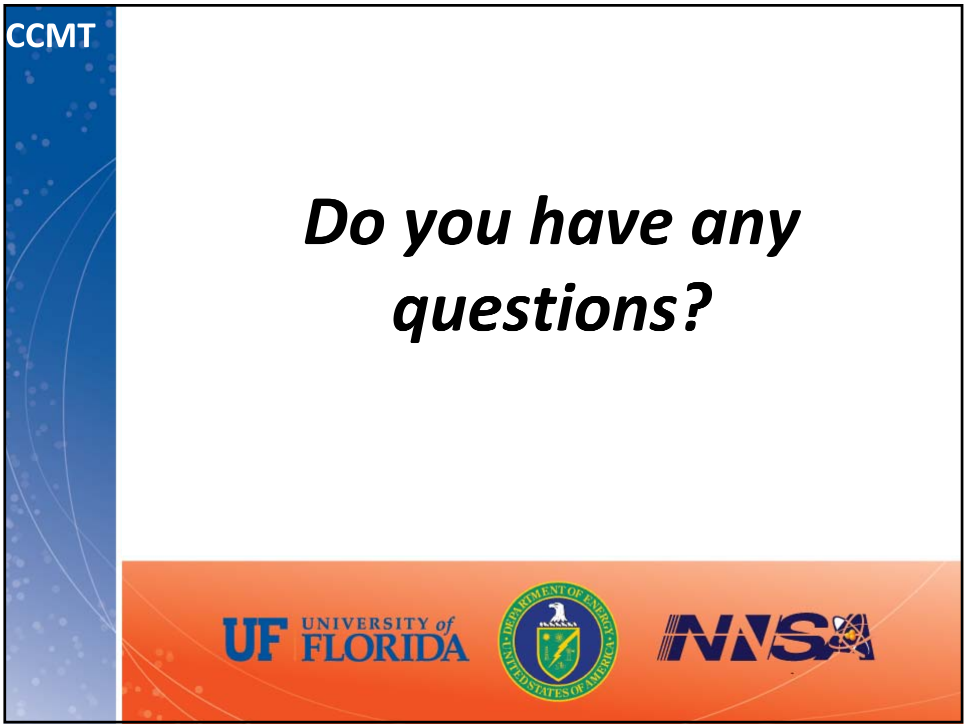

CCMT

Do you have any questions?

CCMT

Full System Integrated Simulations

Bertrand Rollin

UF UNIVERSITY of FLORIDA

NNSA

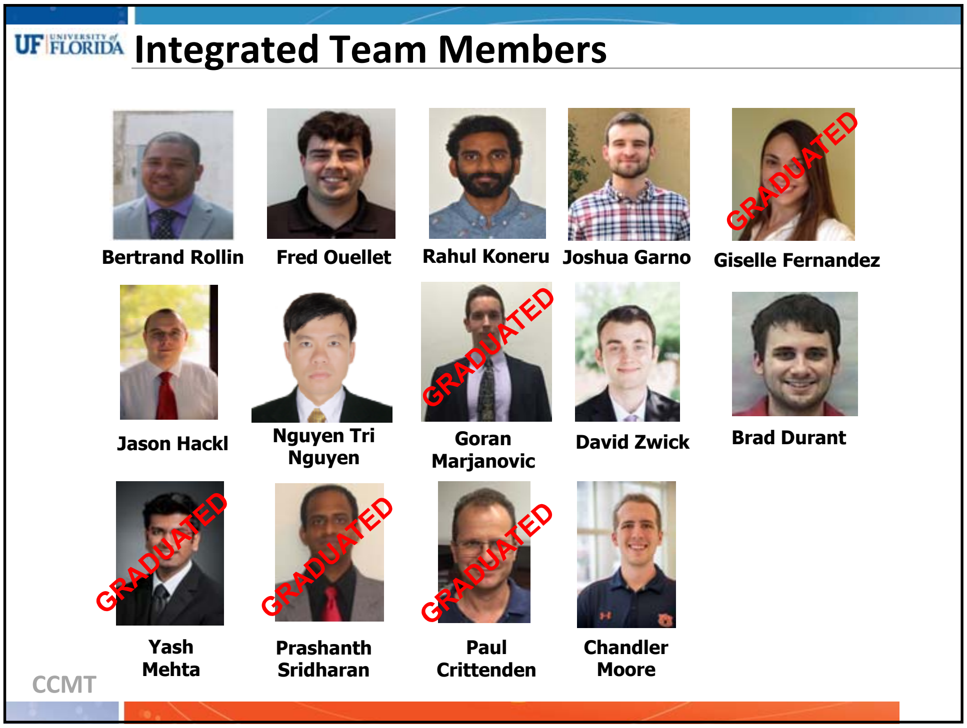

UF UNIVERSITY of FLORIDA

Integrated Team Members

Bertrand Rollin

Fred Ouellet

Rahul Koneru

Joshua Garno

Giselle Fernandez

Jason Hackl

Nguyen Tri Nguyen

Goran Marjanovic

David Zwick

Brad Durant

Yash Mehta

Prashanth Sridharan

Paul Crittenden

Chandler Moore

CCMT

From Rocflu to CMT-nek

	Rocflu	CMT-nek
Method	Finite volume	Discontinuous Galerkin
Mesh	Mixed cells (tetrahedra, hex)	Spectral elements (hex)
Shock capturing	AUSM+up	Artificial viscosity
Time marching	RK3, RK4	TVDRK3
Point particle modeling	2-way coupled	4-way coupled
Parallel capability	O(10^4) MPI ranks, No load-balancing	O(10^6) MPI ranks, load-balanced

- Jason Hackl and David Zwick

CCMT

Motivation

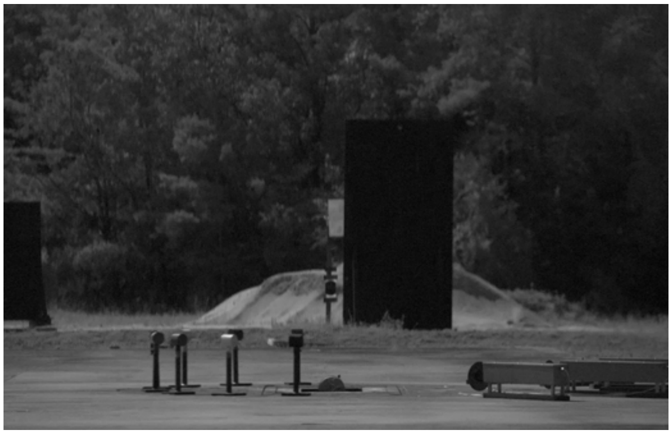

Dispersal of a layer of particles from the detonation of a Comp-B charge during the blast pad experiments performed at AFRL (2017)

CCMT

Sequence of Events and Physical Modeling

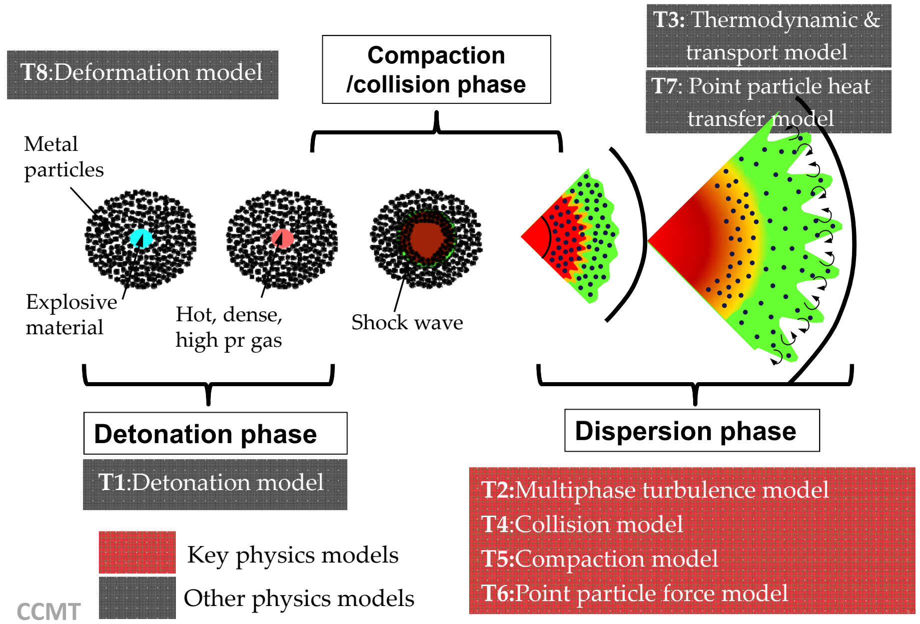

T8:Deformation model

Compaction /collision phase

T3: Thermodynamic & transport model

T7: Point particle heat transfer model

Detonation phase

T1:Detonation model

Dispersion phase

T2:Multiphase turbulence model
T4:Collision model
T5:Compaction model
T6:Point particle force model

🟥 Key physics models
⬛ Other physics models

---

4 Micro/Meso Campaigns & Target Models

- Sandia shock-tube
  - T6: Force coupling and T4: Particle-particle interaction
- ASU expansion fan
  - T2: Multiphase turbulence and T4: Particle-particle interaction
- Eglin microscale
  - T6: Force coupling
- Eglin mesoscale gas-gun
  - T5: Compaction

⬇

- Demonstration problem
  - Yearly hero run

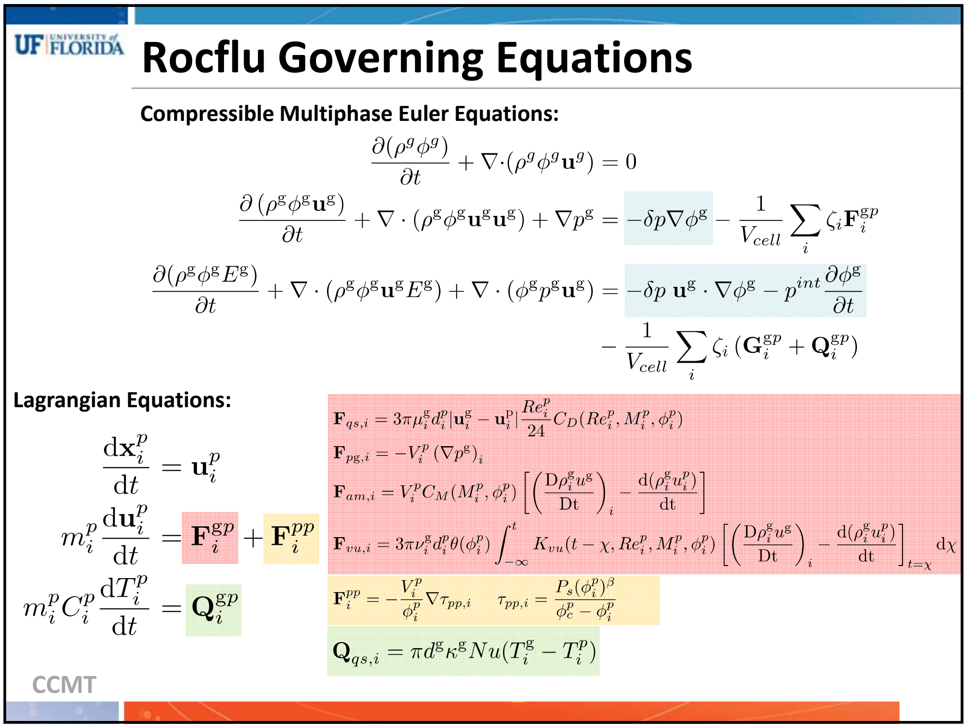

Rocflu Governing Equations

Compressible Multiphase Euler Equations:

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

$$\frac{\partial\left(\rho^{\mathrm{g}}\phi^{\mathrm{g}}\mathbf{u}^{\mathrm{g}}\right)}{\partial t}+\nabla\cdot\left(\rho^{\mathrm{g}}\phi^{\mathrm{g}}\mathbf{u}^{\mathrm{g}}\mathbf{u}^{\mathrm{g}}\right)+\nabla p^{\mathrm{g}}=-\delta p\nabla\phi^{\mathrm{g}}-\frac{1}{V_{cell}}\sum_i\zeta_i\mathbf{F}_i^{\mathrm{g}p}$$

$$\frac{\partial(\rho^{\mathrm{g}}\phi^{\mathrm{g}}E^{\mathrm{g}})}{\partial t}+\nabla\cdot\left(\rho^{\mathrm{g}}\phi^{\mathrm{g}}\mathbf{u}^{\mathrm{g}}E^{\mathrm{g}}\right)+\nabla\cdot\left(\phi^{\mathrm{g}}p^{\mathrm{g}}\mathbf{u}^{\mathrm{g}}\right)=-\delta p\,\mathbf{u}^{\mathrm{g}}\cdot\nabla\phi^{\mathrm{g}}-p^{int}\frac{\partial\phi^{\mathrm{g}}}{\partial t}$$

$$-\frac{1}{V_{cell}}\sum_i\zeta_i\left(\mathbf{G}_i^{\mathrm{g}p}+\mathbf{Q}_i^{\mathrm{g}p}\right)$$

Lagrangian Equations:

$$\frac{\mathrm{d}\mathbf{x}_i^p}{\mathrm{d}t}=\mathbf{u}_i^p$$

$$m_i^p\frac{\mathrm{d}\mathbf{u}_i^p}{\mathrm{d}t}=\mathbf{F}_i^{\mathrm{g}p}+\mathbf{F}_i^{pp}$$

$$m_i^pC_i^p\frac{\mathrm{d}T_i^p}{\mathrm{d}t}=\mathbf{Q}_i^{\mathrm{g}p}$$

$$\mathbf{F}_{qs,i}=3\pi\mu_i^{\mathrm{g}}d_i^p|\mathbf{u}_i^{\mathrm{g}}-\mathbf{u}_i^{\mathrm{p}}|\frac{Re_i^p}{24}C_D(Re_i^p,M_i^p,\phi_i^p)$$

$$\mathbf{F}_{p\mathrm{g},i}=-V_i^p\left(\nabla p^{\mathrm{g}}\right)_i$$

$$\mathbf{F}_{am,i}=V_i^pC_M(M_i^p,\phi_i^p)\left[\left(\frac{\mathrm{D}\rho_i^{\mathrm{g}}u^{\mathrm{g}}}{\mathrm{Dt}}\right)_i-\frac{\mathrm{d}(\rho_i^{\mathrm{g}}u_i^p)}{\mathrm{dt}}\right]$$

$$\mathbf{F}_{vu,i}=3\pi\nu_i^{\mathrm{g}}d_i^p\theta(\phi_i^p)\int_{-\infty}^{t}K_{vu}(t-\chi,Re_i^p,M_i^p,\phi_i^p)\left[\left(\frac{\mathrm{D}\rho_i^{\mathrm{g}}u^{\mathrm{g}}}{\mathrm{Dt}}\right)_i-\frac{\mathrm{d}(\rho_i^{\mathrm{g}}u_i^p)}{\mathrm{dt}}\right]_{t=\chi}\mathrm{d}\chi$$

$$\mathbf{F}_i^{pp}=-\frac{V_i^p}{\phi_i^p}\nabla\tau_{pp,i}\qquad\tau_{pp,i}=\frac{P_s(\phi_i^p)^\beta}{\phi_c^p-\phi_i^p}$$

$$\mathbf{Q}_{qs,i}=\pi d^{\mathrm{g}}\kappa^{\mathrm{g}}Nu(T_i^{\mathrm{g}}-T_i^p)$$

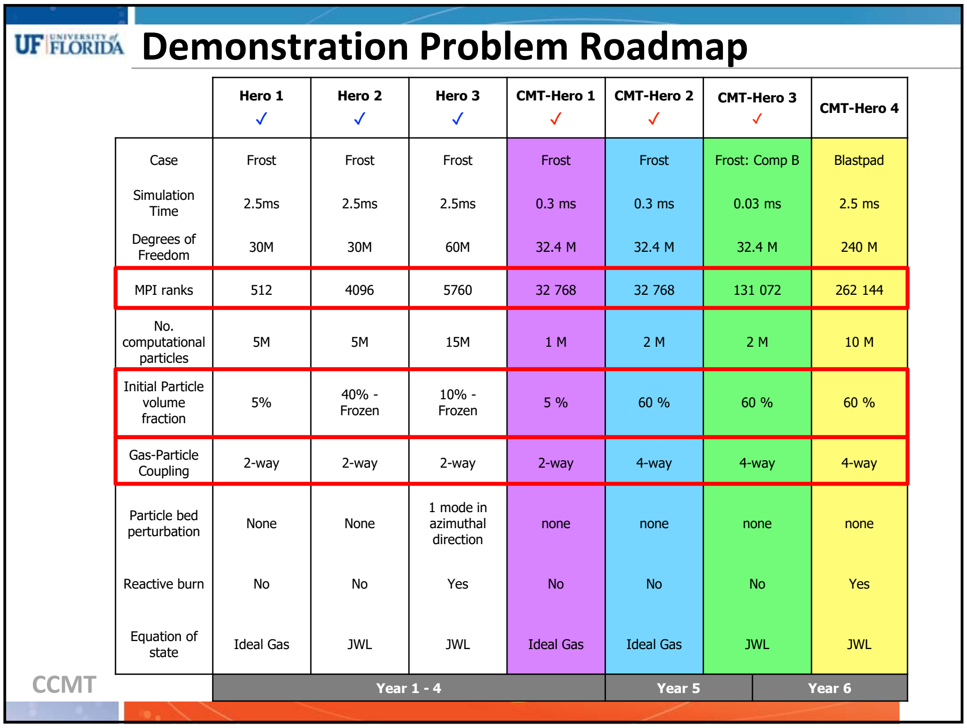

Demonstration Problem Roadmap

	Hero 1 ✓	Hero 2 ✓	Hero 3 ✓	CMT-Hero 1 ✓	CMT-Hero 2 ✓	CMT-Hero 3 ✓	CMT-Hero 4
Case	Frost	Frost	Frost	Frost	Frost	Frost: Comp B	Blastpad
Simulation Time	2.5ms	2.5ms	2.5ms	0.3 ms	0.3 ms	0.03 ms	2.5 ms
Degrees of Freedom	30M	30M	60M	32.4 M	32.4 M	32.4 M	240 M
MPI ranks	512	4096	5760	32 768	32 768	131 072	262 144
No. computational particles	5M	5M	15M	1 M	2 M	2 M	10 M
Initial Particle volume fraction	5%	40% - Frozen	10% - Frozen	5 %	60 %	60 %	60 %
Gas-Particle Coupling	2-way	2-way	2-way	2-way	4-way	4-way	4-way
Particle bed perturbation	None	None	1 mode in azimuthal direction	none	none	none	none
Reactive burn	No	No	Yes	No	No	No	Yes
Equation of state	Ideal Gas	JWL	JWL	Ideal Gas	Ideal Gas	JWL	JWL
	Year 1 - 4				Year 5	Year 6	

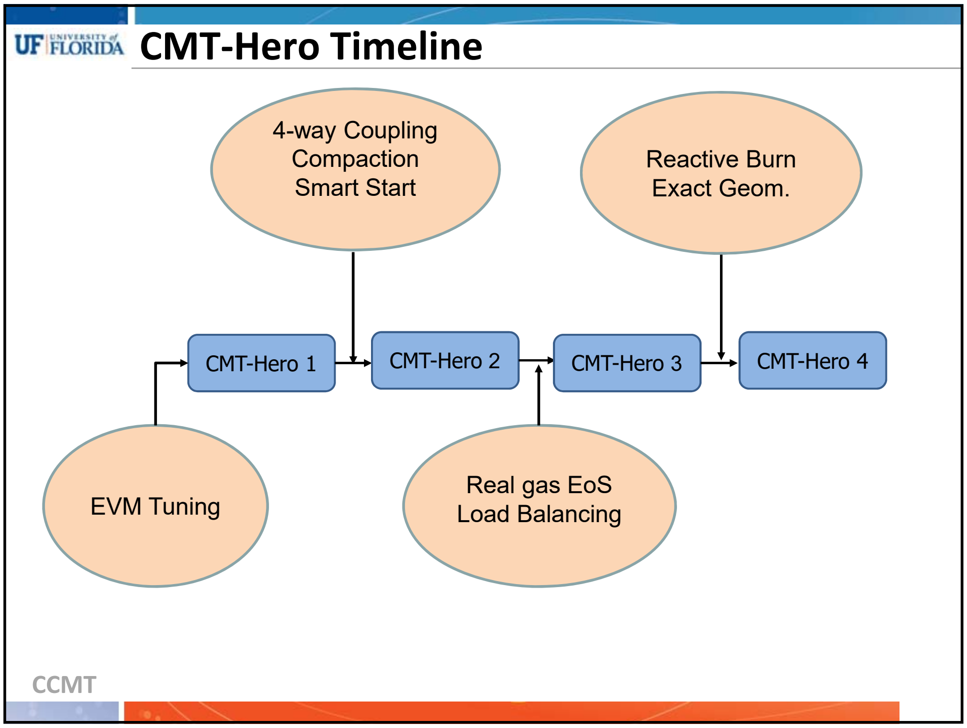

CMT-Hero Timeline

4-way Coupling
Compaction
Smart Start

Reactive Burn
Exact Geom.

CMT-Hero 1

CMT-Hero 2

CMT-Hero 3

CMT-Hero 4

EVM Tuning

Real gas EoS
Load Balancing

CCMT

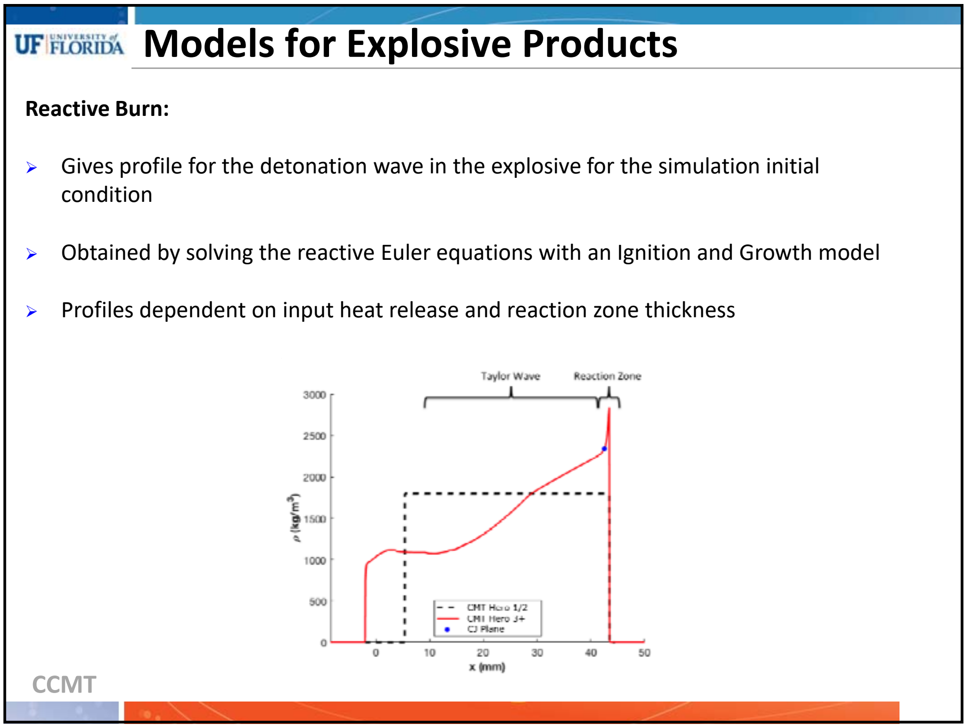

Models for Explosive Products

Reactive Burn:

- Gives profile for the detonation wave in the explosive for the simulation initial condition
- Obtained by solving the reactive Euler equations with an Ignition and Growth model
- Profiles dependent on input heat release and reaction zone thickness

CCMT

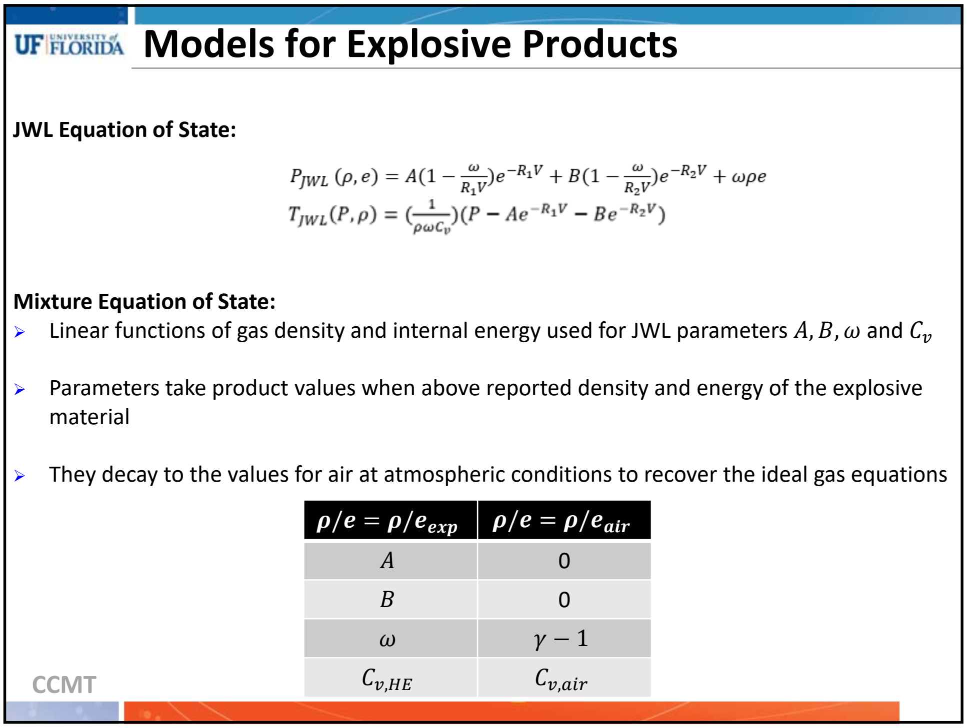

Models for Explosive Products

JWL Equation of State:

$$P_{JWL}(\rho,e) = A(1-\frac{\omega}{R_1 V})e^{-R_1 V} + B(1-\frac{\omega}{R_2 V})e^{-R_2 V} + \omega\rho e$$

$$T_{JWL}(P,\rho) = (\frac{1}{\rho\omega C_v})(P - Ae^{-R_1 V} - Be^{-R_2 V})$$

Mixture Equation of State:

- Linear functions of gas density and internal energy used for JWL parameters A, B, ω and C_v
- Parameters take product values when above reported density and energy of the explosive material
- They decay to the values for air at atmospheric conditions to recover the ideal gas equations

$\rho/e = \rho/e_{exp}$	$\rho/e = \rho/e_{air}$
A	0
B	0
ω	$\gamma - 1$
$C_{v,HE}$	$C_{v,air}$

CCMT

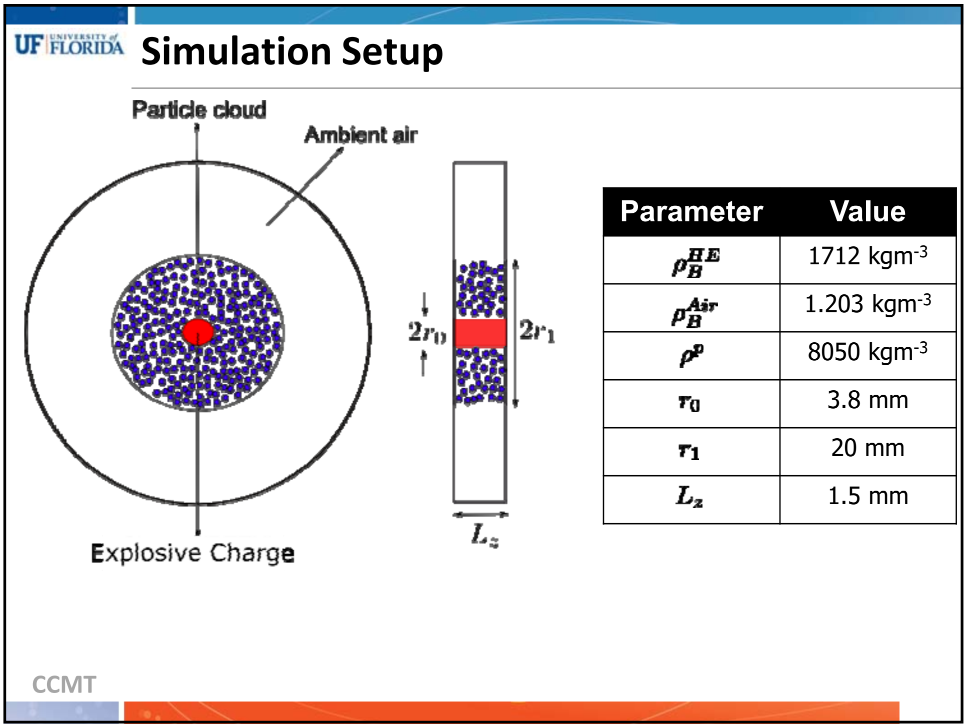

Simulation Setup

Parameter	Value
ρ_B^{HE}	1712 kgm^{-3}
ρ_B^{Air}	1.203 kgm^{-3}
ρ^p	8050 kgm^{-3}
r_0	3.8 mm
r_1	20 mm
L_z	1.5 mm

CCMT

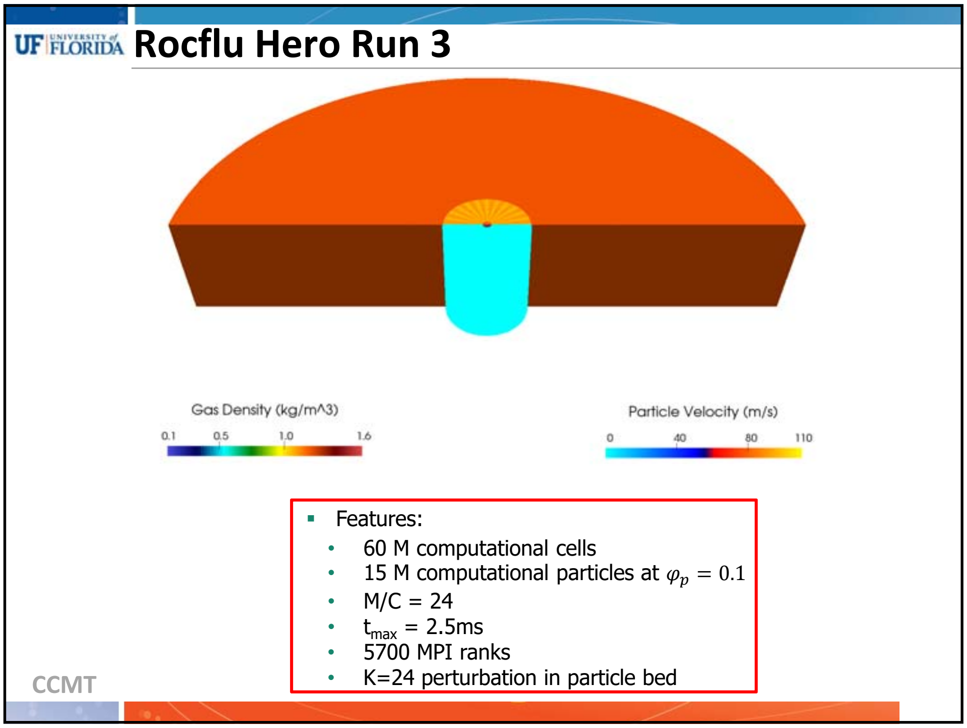

Rocflu Hero Run 3

- Features:
  - 60 M computational cells
  - 15 M computational particles at $\varphi_p = 0.1$
  - M/C = 24
  - t_{max} = 2.5ms
  - 5700 MPI ranks
  - K=24 perturbation in particle bed

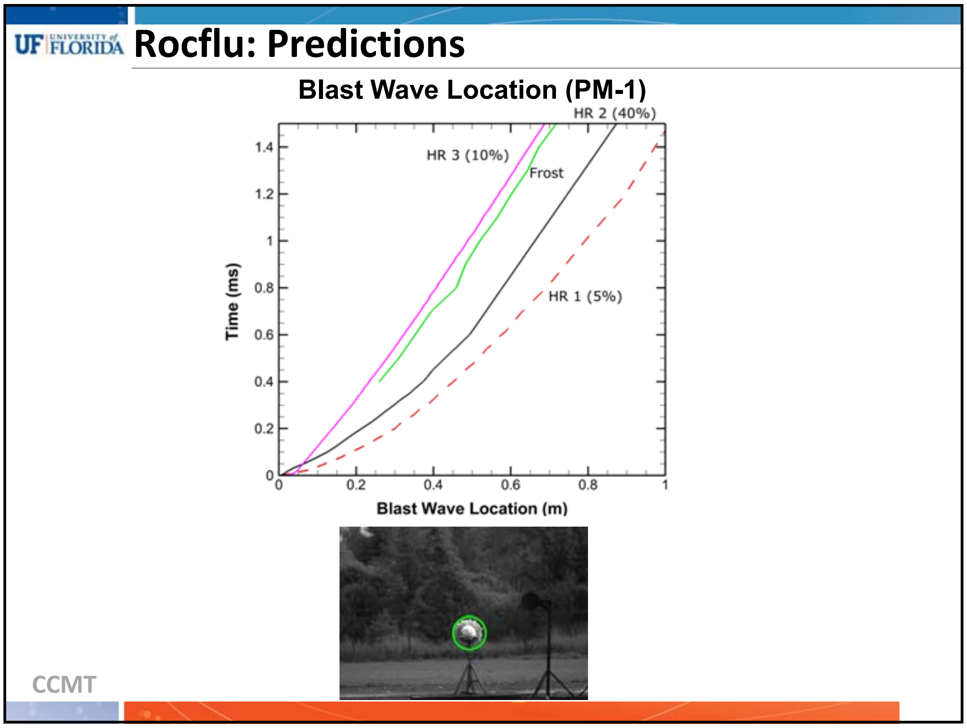

Rocflu: Predictions

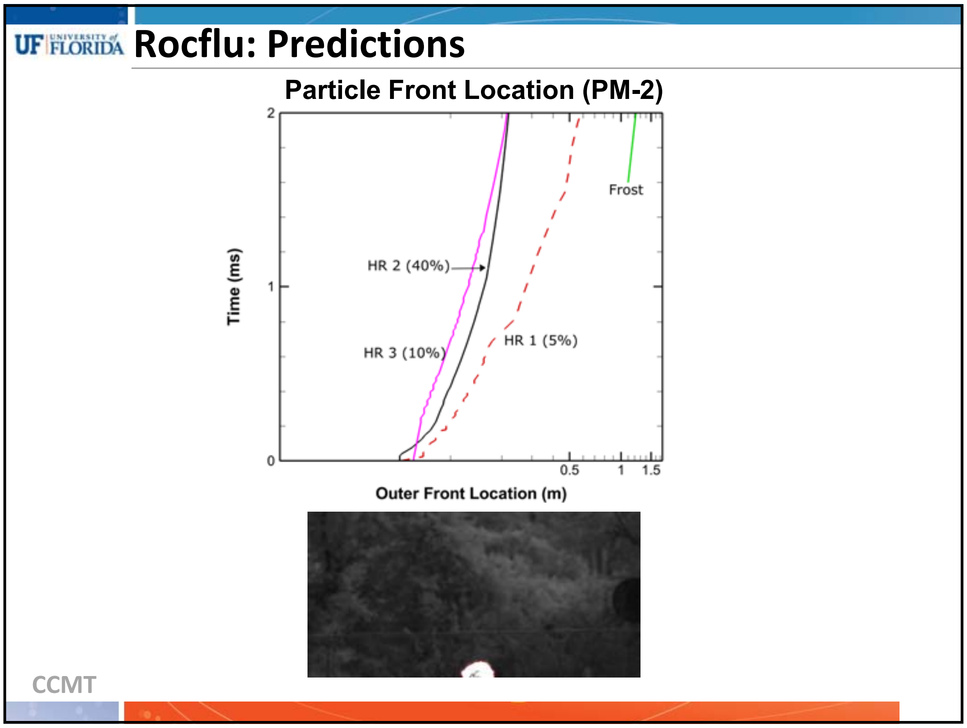

UF | UNIVERSITY of FLORIDA

Rocflu: Predictions

Particle Front Location (PM-2)

CCMT

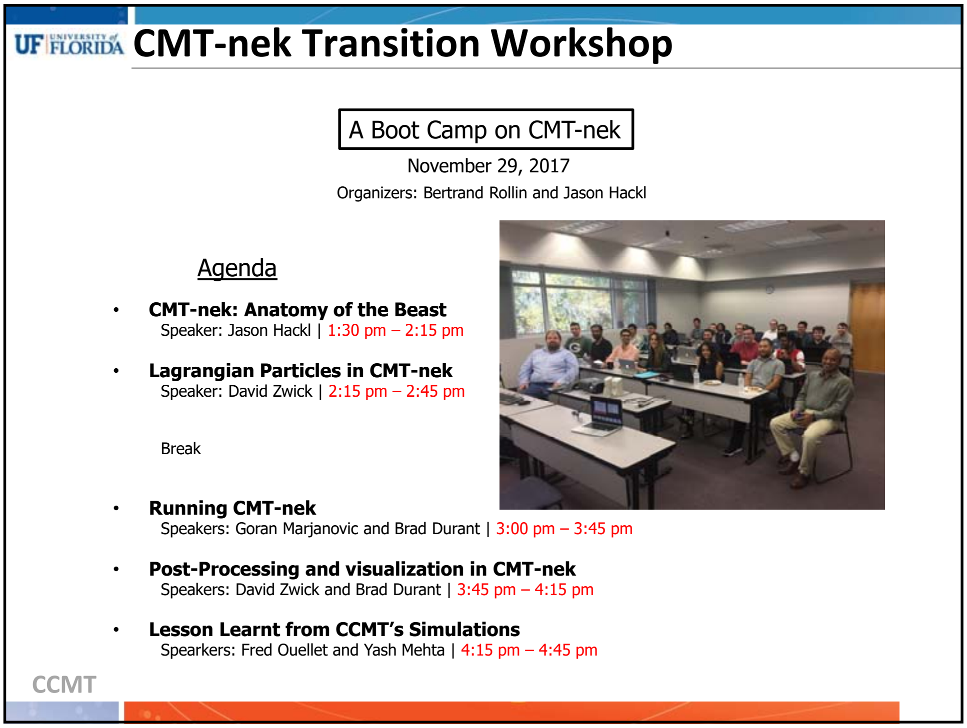

UF | UNIVERSITY of FLORIDA

CMT-nek Transition Workshop

A Boot Camp on CMT-nek

November 29, 2017

Organizers: Bertrand Rollin and Jason Hackl

Agenda

- **CMT-nek: Anatomy of the Beast**
  Speaker: Jason Hackl | 1:30 pm – 2:15 pm
- **Lagrangian Particles in CMT-nek**
  Speaker: David Zwick | 2:15 pm – 2:45 pm

  Break

- **Running CMT-nek**
  Speakers: Goran Marjanovic and Brad Durant | 3:00 pm – 3:45 pm
- **Post-Processing and visualization in CMT-nek**
  Speakers: David Zwick and Brad Durant | 3:45 pm – 4:15 pm
- **Lesson Learnt from CCMT's Simulations**
  Spearkers: Fred Ouellet and Yash Mehta | 4:15 pm – 4:45 pm

CCMT

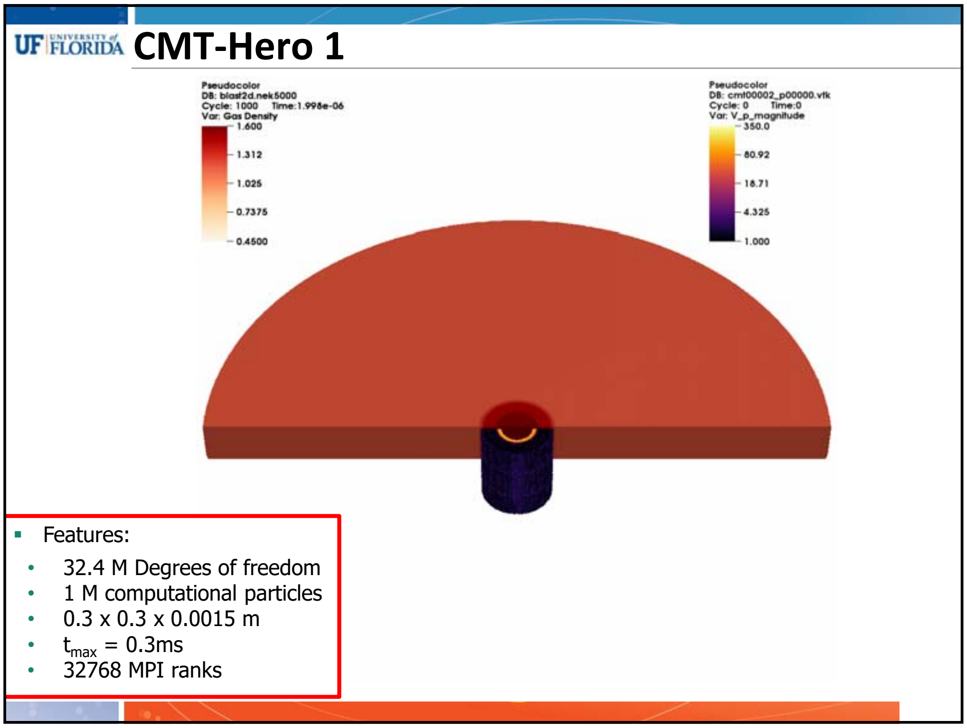

CMT-Hero 1

- Features:
  - 32.4 M Degrees of freedom
  - 1 M computational particles
  - 0.3 x 0.3 x 0.0015 m
  - t_{max} = 0.3ms
  - 32768 MPI ranks

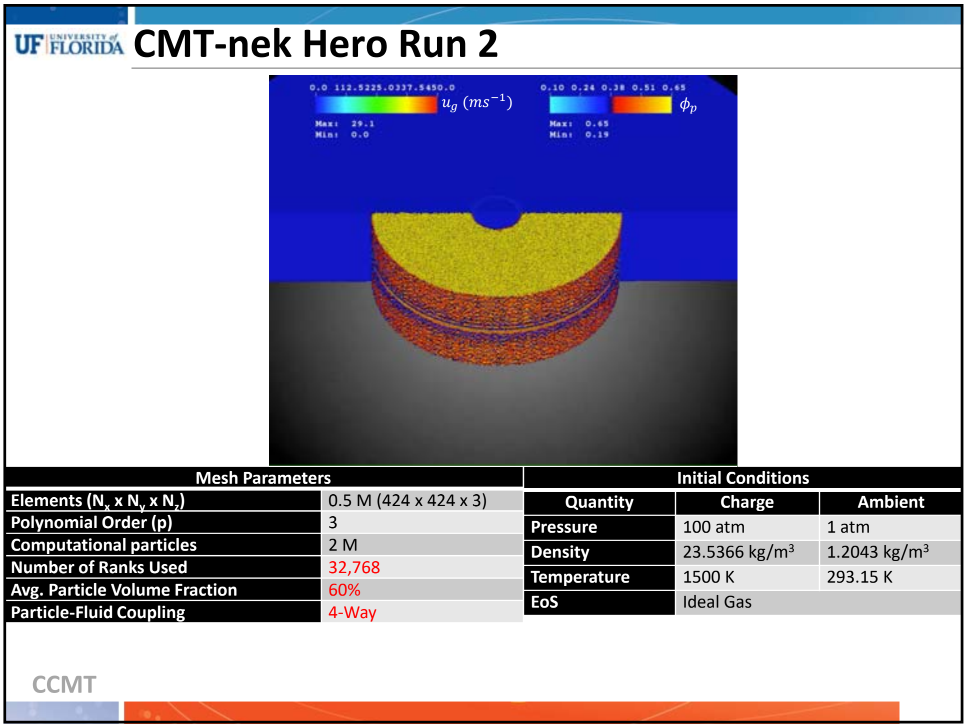

CMT-nek Hero Run 2

Mesh Parameters	
Elements (N_x x N_y x N_z)	0.5 M (424 x 424 x 3)
Polynomial Order (p)	3
Computational particles	2 M
Number of Ranks Used	32,768
Avg. Particle Volume Fraction	60%
Particle-Fluid Coupling	4-Way

Initial Conditions		
Quantity	Charge	Ambient
Pressure	100 atm	1 atm
Density	23.5366 kg/m^3	1.2043 kg/m^3
Temperature	1500 K	293.15 K
EoS	Ideal Gas	

CCMT

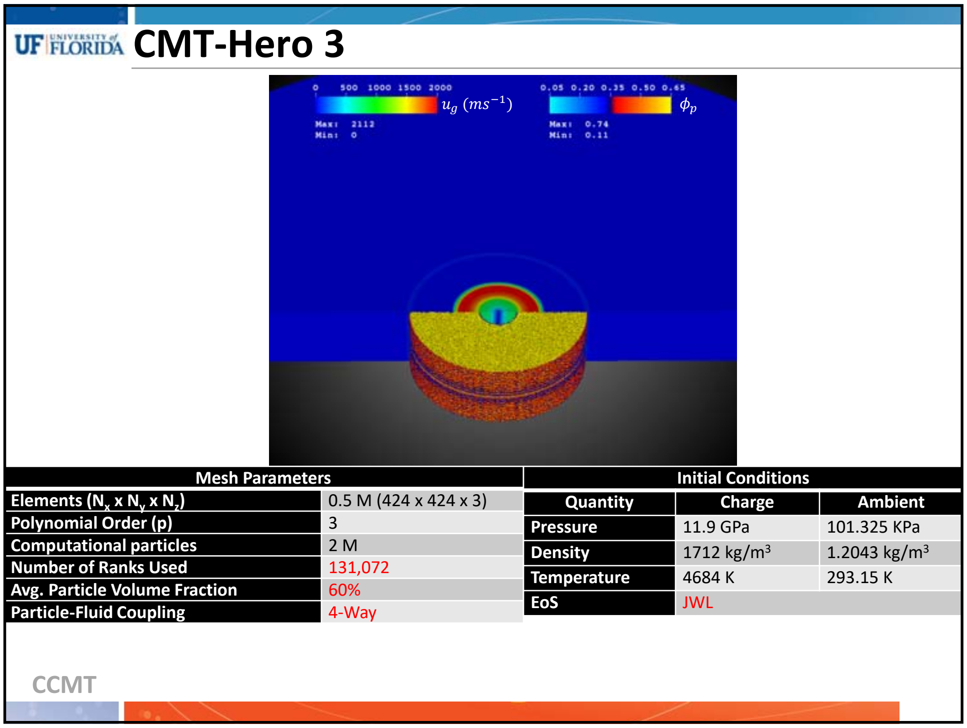

CMT-Hero 3

Mesh Parameters	
Elements (N_x x N_y x N_z)	0.5 M (424 x 424 x 3)
Polynomial Order (p)	3
Computational particles	2 M
Number of Ranks Used	131,072
Avg. Particle Volume Fraction	60%
Particle-Fluid Coupling	4-Way

Initial Conditions		
Quantity	**Charge**	**Ambient**
Pressure	11.9 GPa	101.325 KPa
Density	1712 kg/m^3	1.2043 kg/m^3
Temperature	4684 K	293.15 K
EoS	JWL	

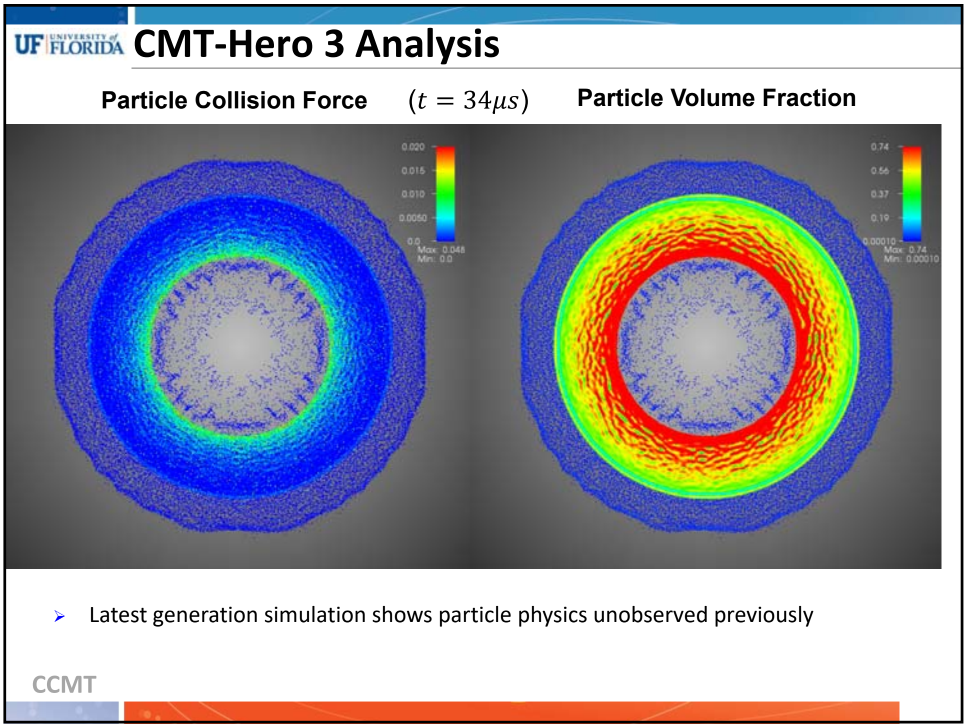

CMT-Hero 3 Analysis

Particle Collision Force $(t = 34\mu s)$ **Particle Volume Fraction**

- Latest generation simulation shows particle physics unobserved previously

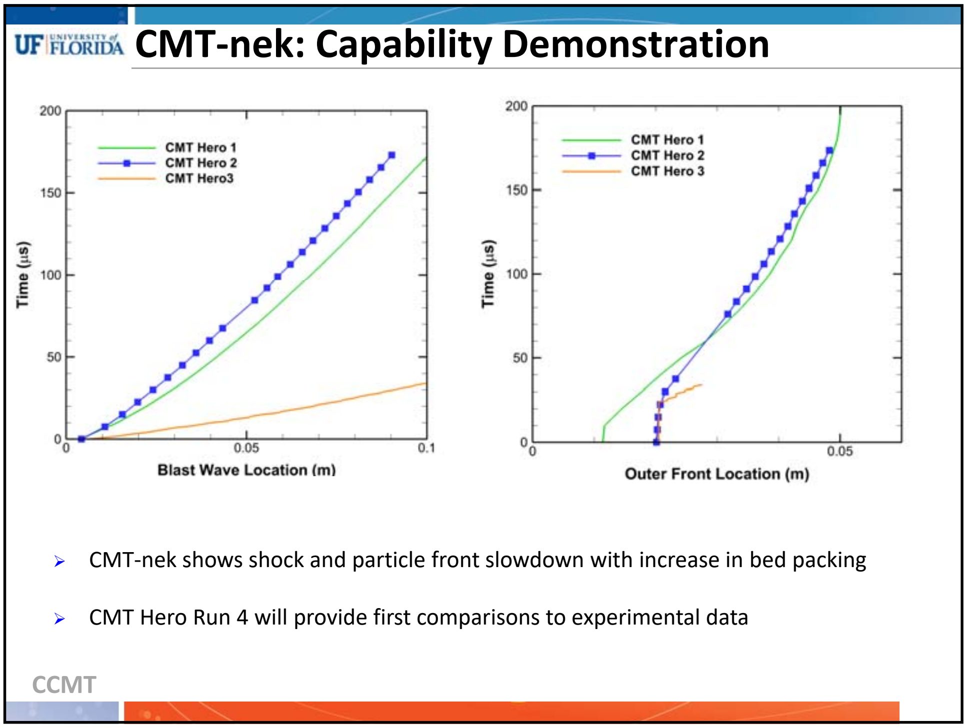

CMT-nek: Capability Demonstration

- CMT-nek shows shock and particle front slowdown with increase in bed packing
- CMT Hero Run 4 will provide first comparisons to experimental data

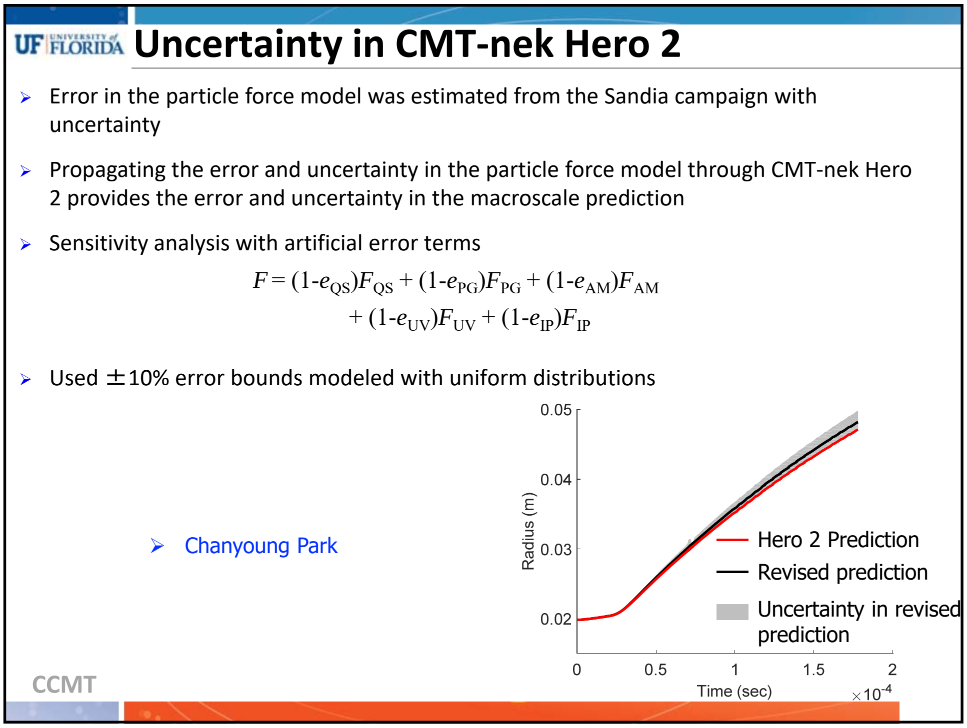

Uncertainty in CMT-nek Hero 2

- Error in the particle force model was estimated from the Sandia campaign with uncertainty
- Propagating the error and uncertainty in the particle force model through CMT-nek Hero 2 provides the error and uncertainty in the macroscale prediction
- Sensitivity analysis with artificial error terms

$$F = (1\text{-}e_{\mathrm{QS}})F_{\mathrm{QS}} + (1\text{-}e_{\mathrm{PG}})F_{\mathrm{PG}} + (1\text{-}e_{\mathrm{AM}})F_{\mathrm{AM}}$$
$$+ (1\text{-}e_{\mathrm{UV}})F_{\mathrm{UV}} + (1\text{-}e_{\mathrm{IP}})F_{\mathrm{IP}}$$

- Used ±10% error bounds modeled with uniform distributions

- Chanyoung Park

The Channeling Instability: Particle Speed

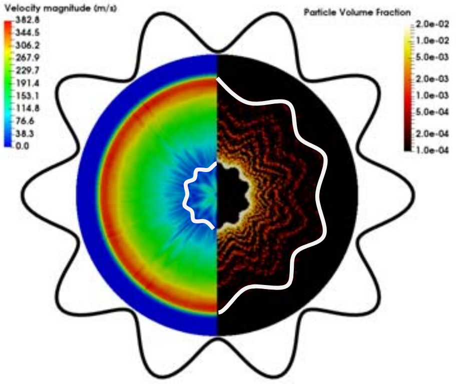

Gas velocity (left) and particle volume fraction (right) contours at final time (t=500μs) for the unimodal case k=10. Black contour represents initial (t=0) unimodal particle volume contours. White contours are schematics of the final (t= 500μs) outer particle front location

- Particles in radial sectors with low particle volume travel faster
- At later times the outer front of particles location flips the initial particle volume profile
- The number of instabilities forming is directly dependent on the initial modal perturbation

The Channeling Instability: Particle Migration

- Animations below show particle volume contours for three different initial perturbations
- The animations show the migration of particles from sectors with lower PV to sectors with higher PV as time goes

$Unimodal\ Perturbation$
$k = 8$

$Trimodal\ Perturbation$
$(A_1, A_2, A_3, k_1, k_2, k_3, \Phi_{12}, \Phi_{13})$
(0.04, 0.12, 0.07, 1, 12, 8, 6.00, 0.38)

Eglin Blastpad Experimental Setup

(Barreto et al. 2015)

- Six shots
  - 2 bare charges
  - 1 charge w/tungsten (n)
  - 2 charges w/steel (n)
  - 1 charge w/steel (u-n)

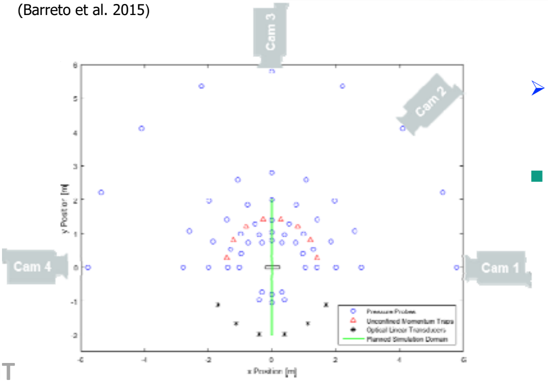

➢ Angela Diggs and Kyle Hughes

- Instrumentation:
  - 54 pressure probes
  - 8 Momentum traps
  - 4 high speed video cameras
  - 6 Linear optical transducers

CCMT

Demonstration Problem – A Summary

- ✓ Last Rocflu Hero Run completed
- ✓ UQ analysis continues
- ✓ Study of Physics of the Demonstration problem continues
- ✓ CMT-Hero 1 completed
- ✓ CMT-Hero 2 completed
- ✓ CMT-Hero 3 completed
- ✓ Blastpad tests in progress

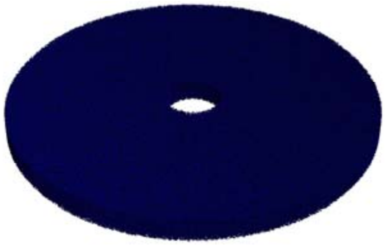

➢ David Zwick

CCMT

4 Micro/Meso Campaigns & Target Models

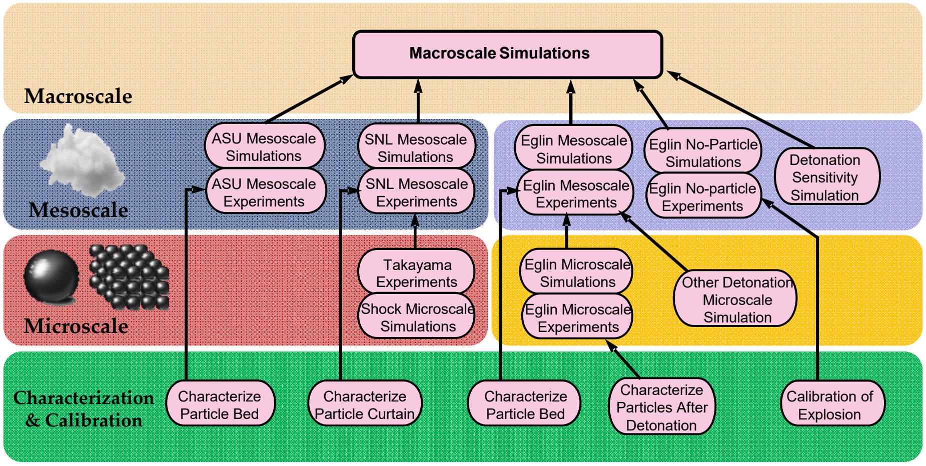

Macroscale Simulations

Macroscale

Mesoscale

ASU Mesoscale Simulations

ASU Mesoscale Experiments

SNL Mesoscale Simulations

SNL Mesoscale Experiments

Eglin Mesoscale Simulations

Eglin Mesoscale Experiments

Eglin No-Particle Simulations

Eglin No-particle Experiments

Detonation Sensitivity Simulation

Microscale

Takayama Experiments

Shock Microscale Simulations

Eglin Microscale Simulations

Eglin Microscale Experiments

Other Detonation Microscale Simulation

Characterization & Calibration

Characterize Particle Bed

Characterize Particle Curtain

Characterize Particle Bed

Characterize Particles After Detonation

Calibration of Explosion

Eglin Microscale/Mesoscale Experiments

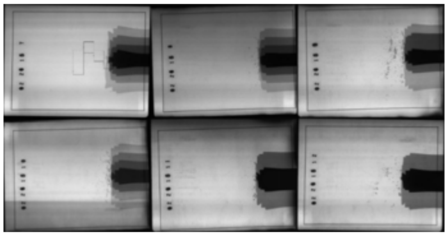

- Eglin Microscale campaign is a joint effort between UQ, experimental and simulation teams
- More details in upcoming talks:
  - Kyle Hughes (Experiments)
  - Chanyoung Park (UQ)
  - Josh Garno (Simulations)

Eglin Microscale Simulations

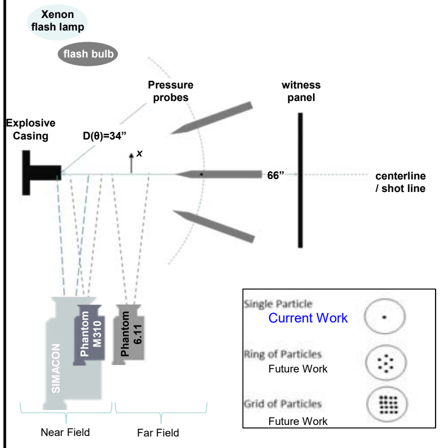

Overhead schematic of the test set-up

Simulation Parameters	
Cells	5 M
Barrel Radius	6.35 mm
Charge length	3.81 cm
Particle-Fluid Coupling	1-Way

Particle Parameters	
Diameter	1.992 to 2.04 mm
Density	15.59 to 16.01 gcc
Vertical position from axis	0, 0.125, 0.254 mm

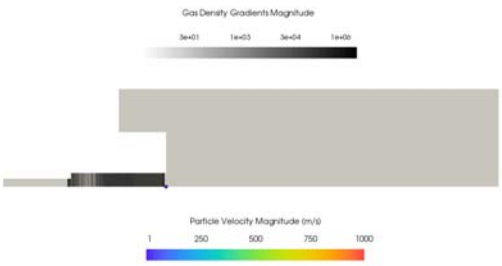

- Josh Garno

Mesoscale Simulation of Mesoscale Experiment

P-Rad Mesoscale Experiment (x-ray)

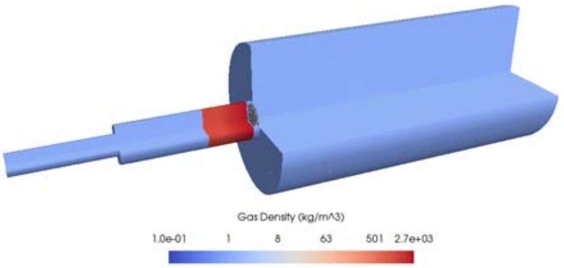

Simulation Details:

- Real gas equation of state (JWL)
- Point-particle model
- Explosive modeled with a reactive burn initial profile

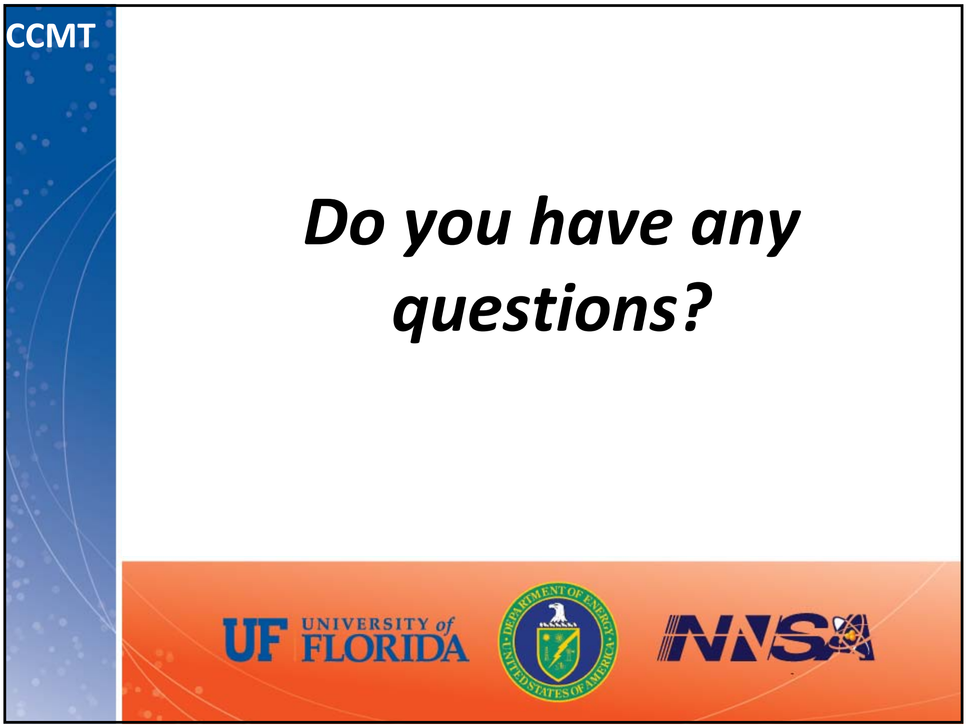

CCMT

Do you have any questions?

CCMT

Experimental Studies of Gas-Particle Mixtures Under Sudden Expansion at ASU

Heather Zunino
PhD Candidate
Dr. Ronald Adrian
Advisor

Problem Statement and Goals

- Experimental multi-phase studies involving compressible flow are complicated
  - Air and solid particles may move separately
  - Particles generate turbulence
- Need for a simple 1D flow experiment that can be used for early validation of the computational codes developed by the PSAAP center.
- Simpler physics involved than the PSAAP capstone experiment
- Perform experiments on existing shock tube setup
- Examine expansion fan, flow structures, turbulence, and instabilities
- Provide data for early-stage validation of computational codes developed by the PSAAP Center

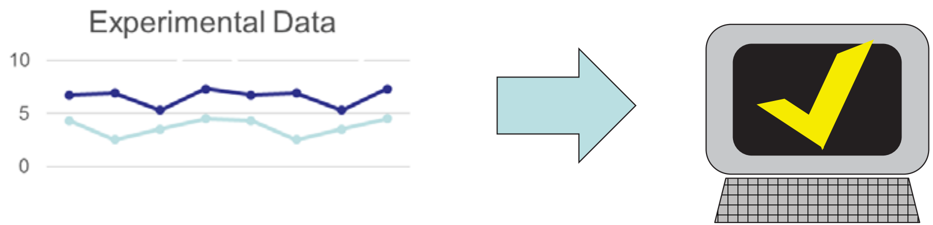

CCMT

Experiment Description

- 1 meter glass tube
    - Cylindrical footprint
    - Inner diameter: 3.9cm
- Particle bed
- Diaphragm
    - Tape
- High-speed camera
- Measurements
    - Gas velocity
    - Particle volume concentration
    - Particle interface
- Parameters: particle size, bed height, and pressure ratio

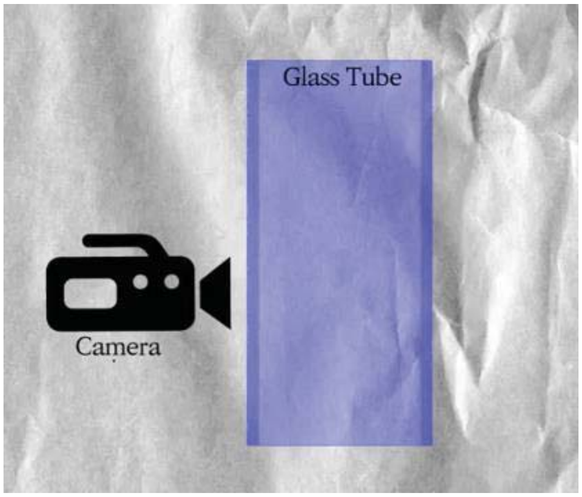

Expansion Waves and Timing Diagram

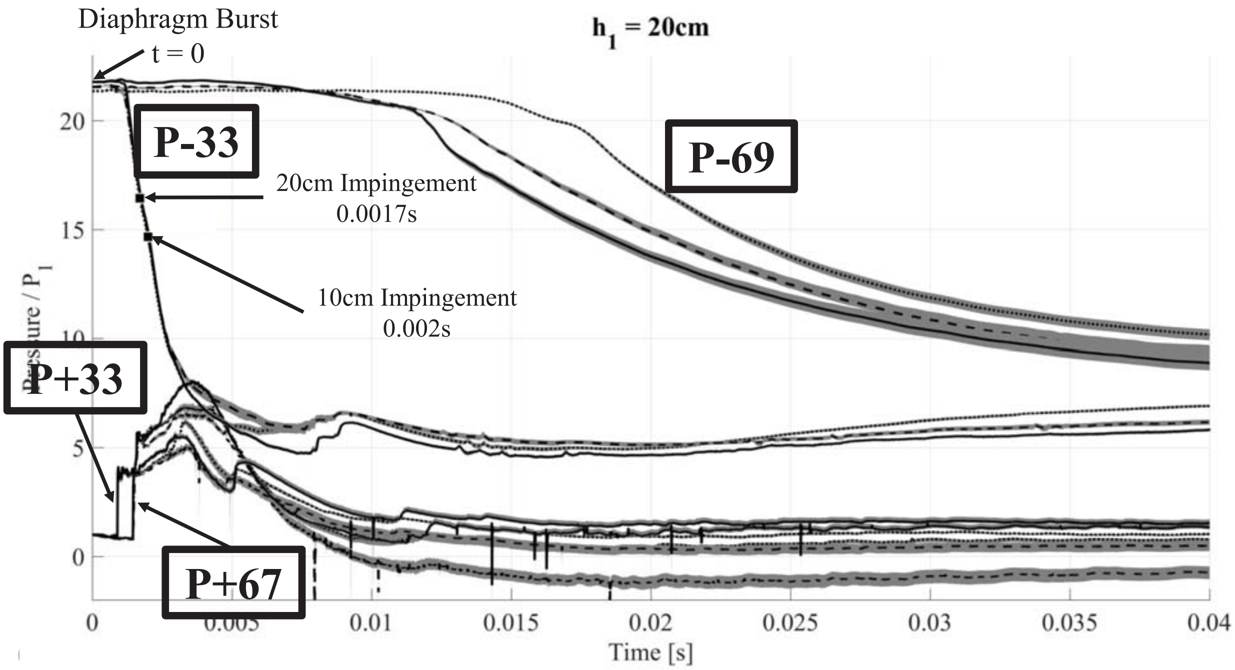

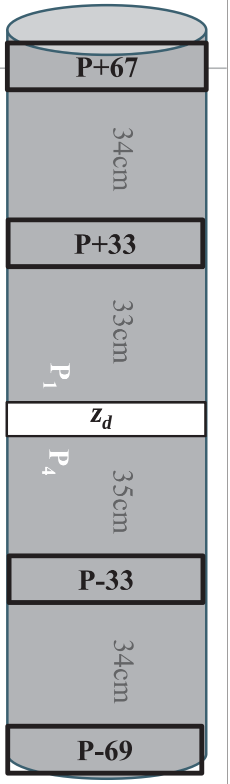

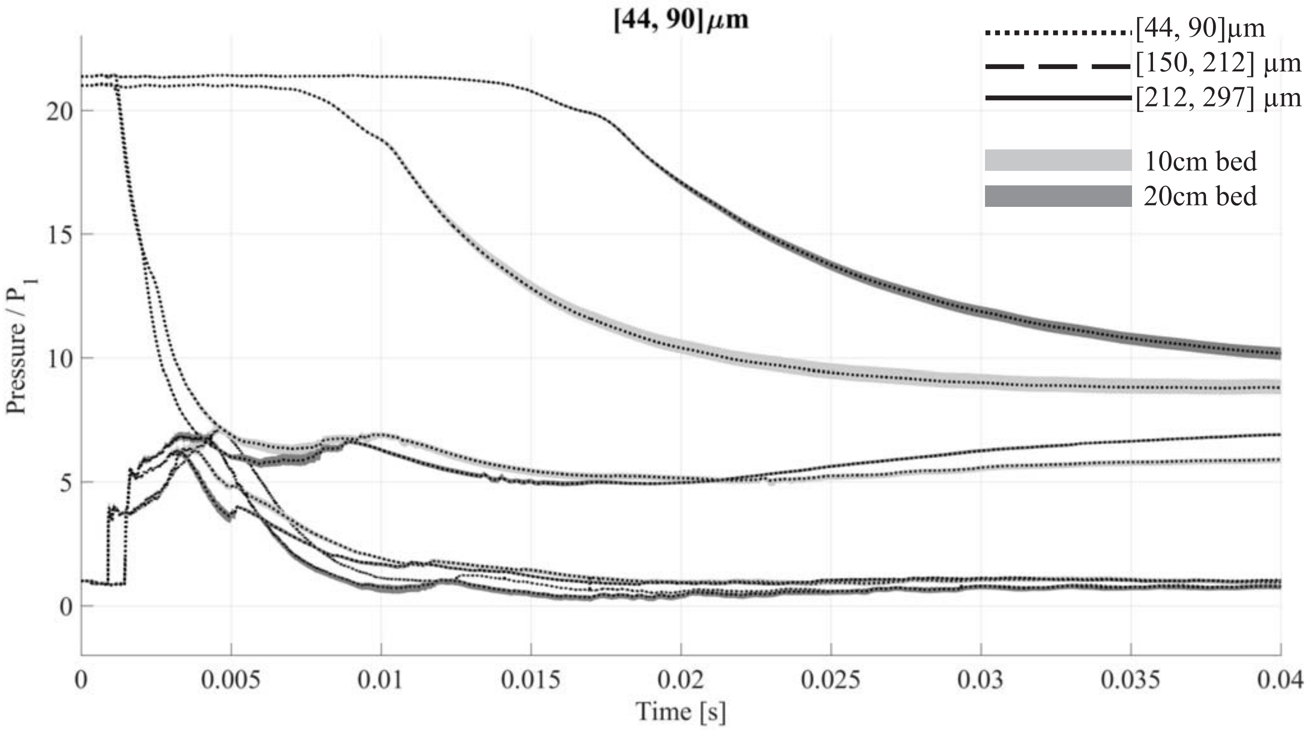
[44, 90]μm
Pressure / P1
Time [s]
0
5
10
15
20
0
0.005
0.01
0.015
0.02
0.025
0.03
0.035
0.04
[44, 90]µm
[150, 212] µm
[212, 297] µm
10cm bed
20cm bed

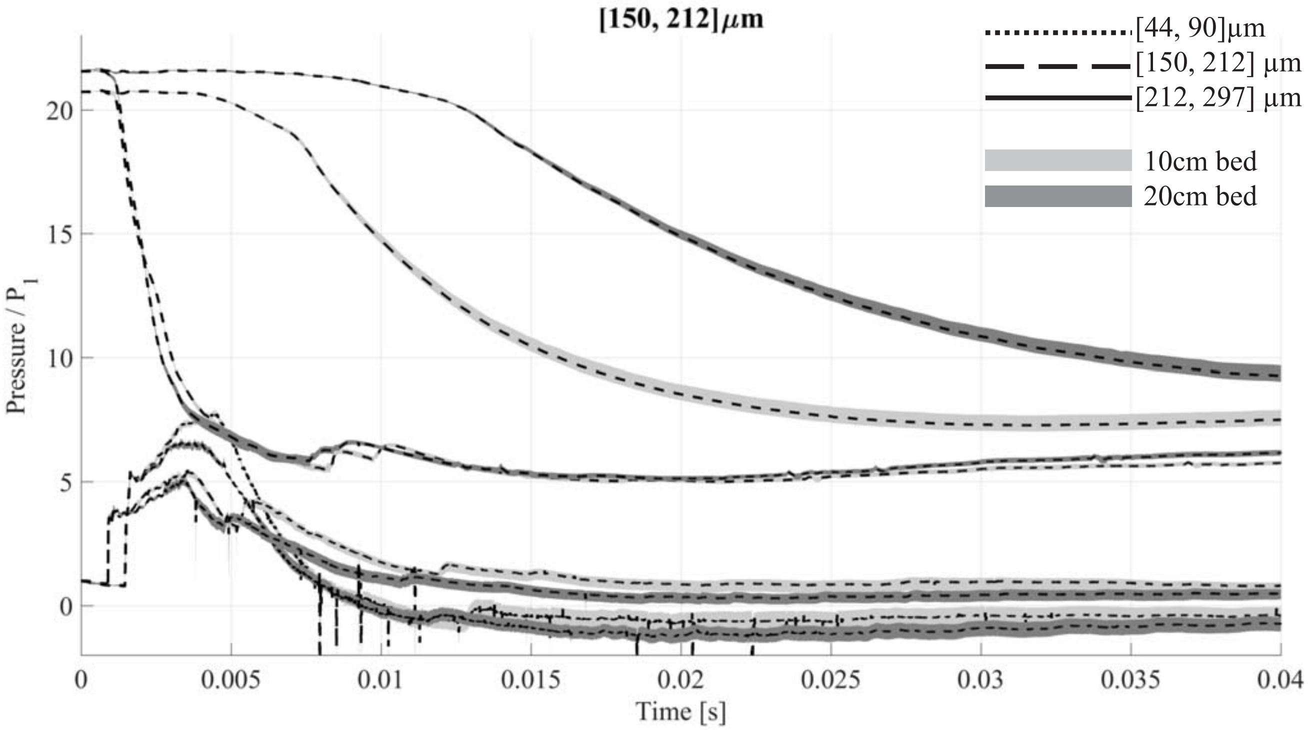
[150, 212]μm
Pressure / P1
Time [s]
0
5
10
15
20
0
0.005
0.01
0.015
0.02
0.025
0.03
0.035
0.04
[44, 90]µm
[150, 212] µm
[212, 297] µm
10cm bed
20cm bed

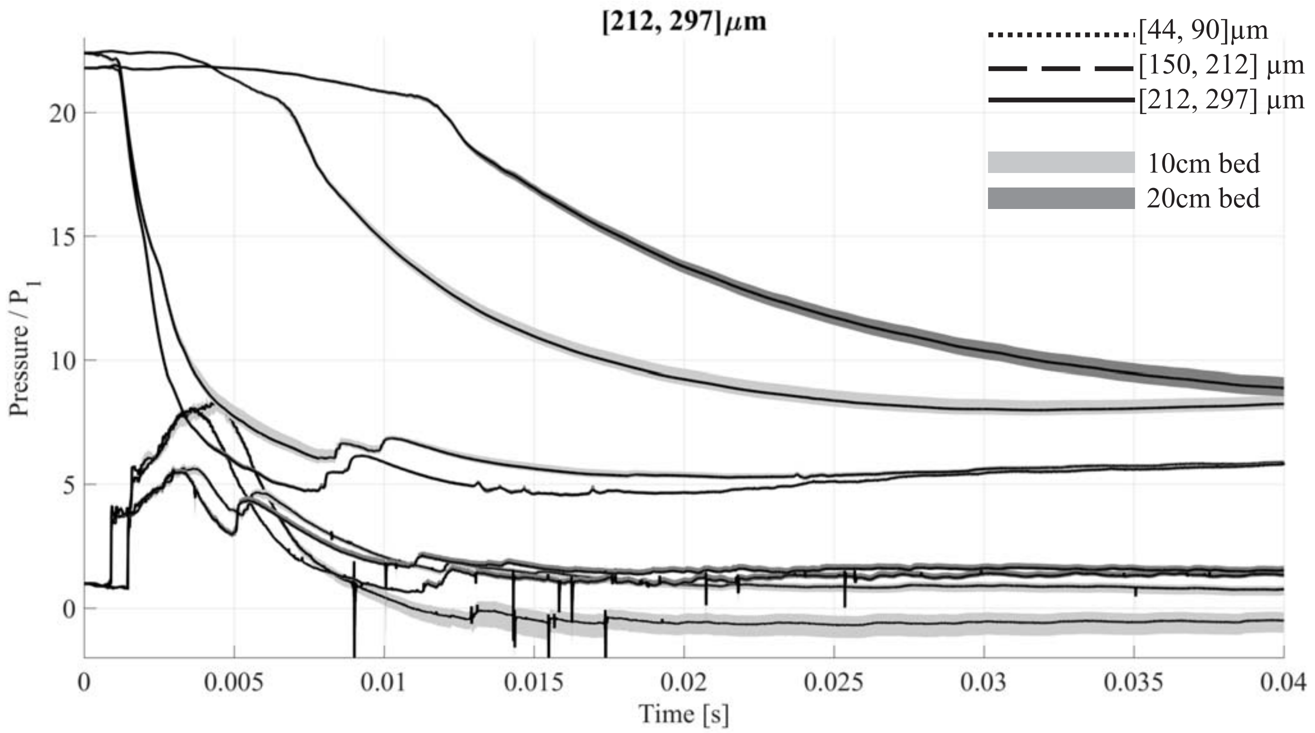
[212, 297]μm
[44, 90]μm
[150, 212] μm
[212, 297] μm
10cm bed
20cm bed
Pressure / P_1
0
5
10
15
20
0
0.005
0.01
0.015
0.02
0.025
0.03
0.035
0.04
Time [s]

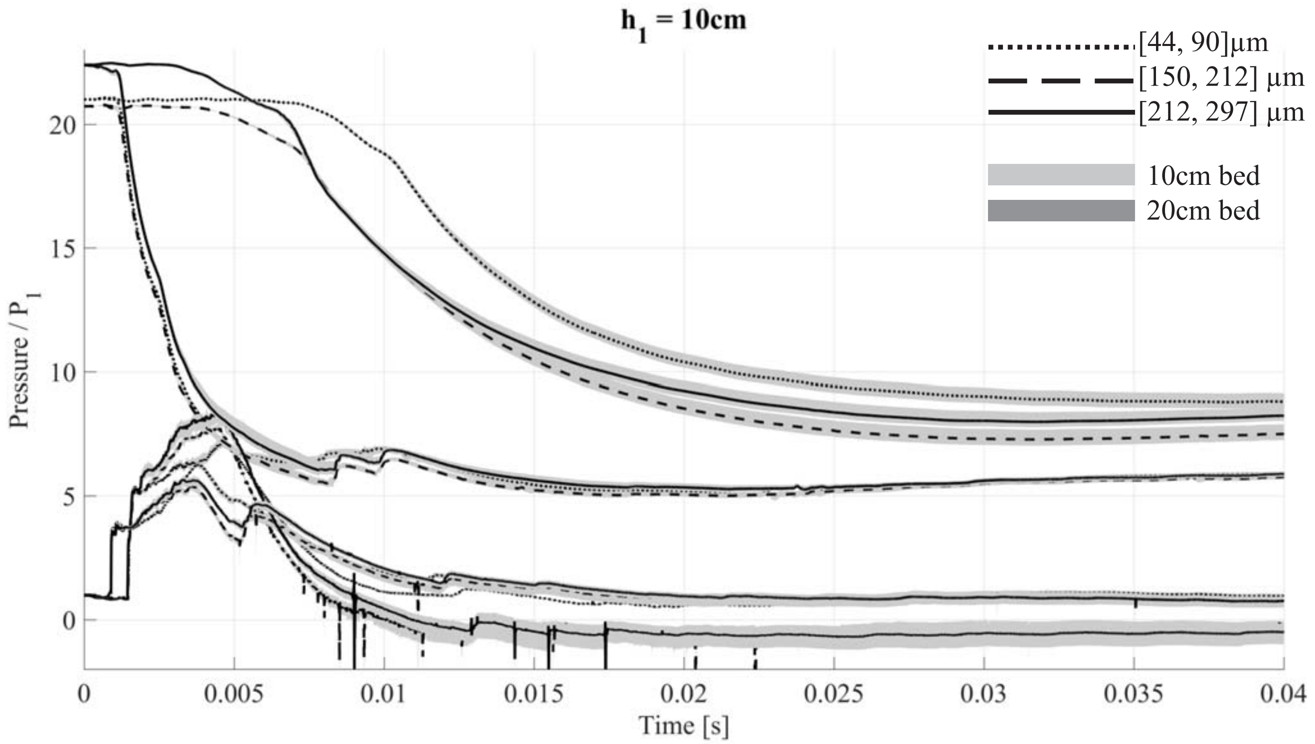
h_1 = 10cm
[44, 90]μm
[150, 212] μm
[212, 297] μm
10cm bed
20cm bed
Pressure / P_1
0
5
10
15
20
0
0.005
0.01
0.015
0.02
0.025
0.03
0.035
0.04
Time [s]

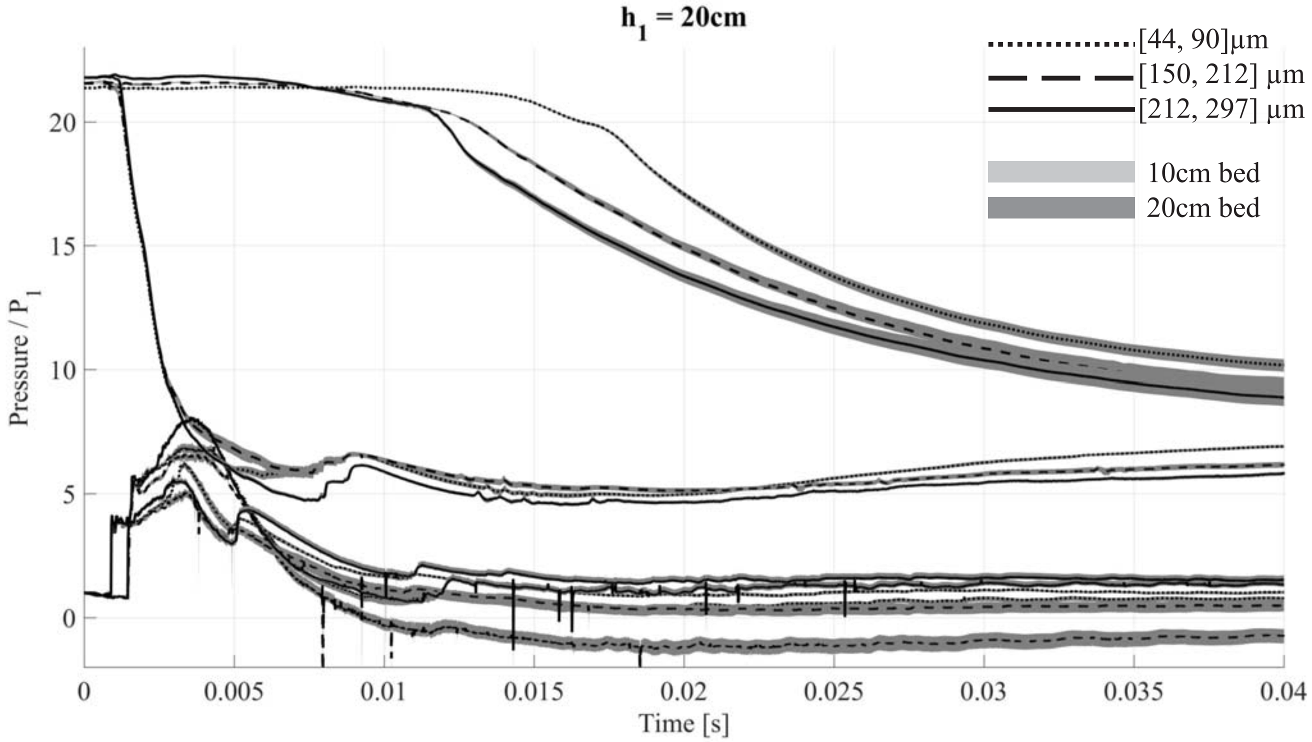

h_1 = 20cm

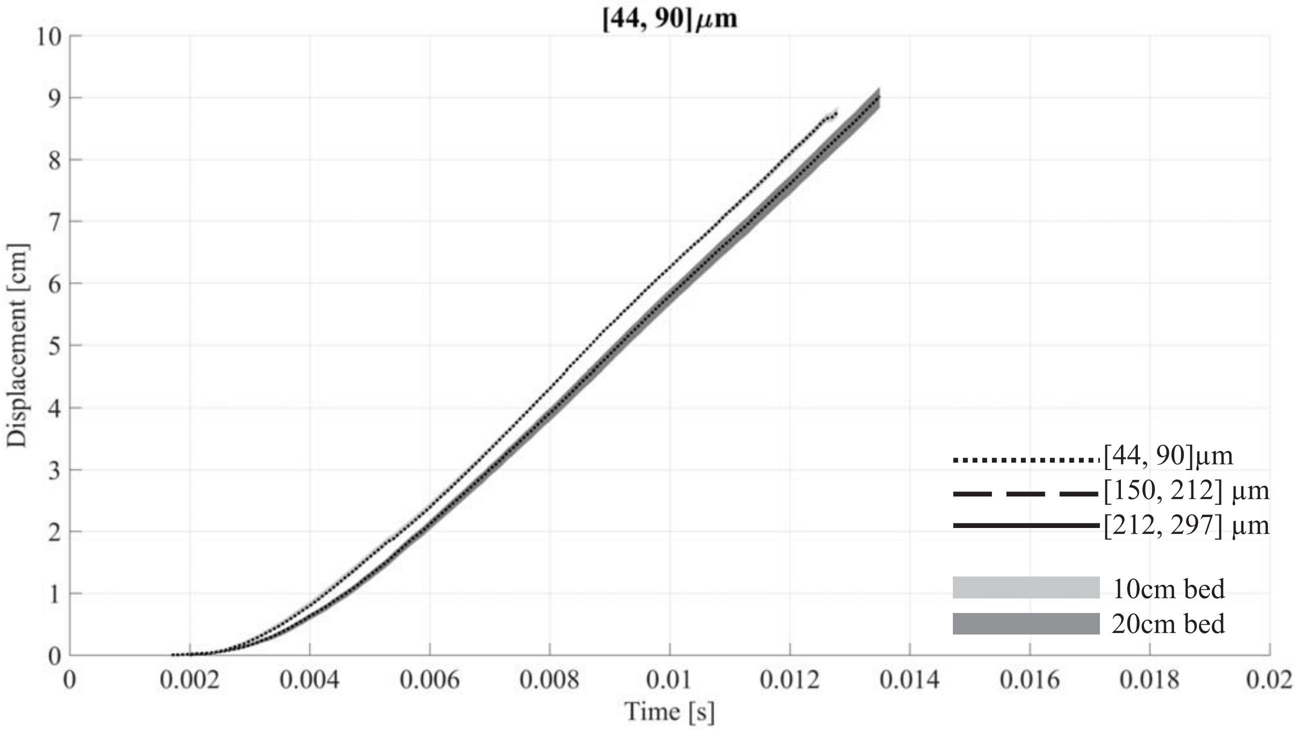

[44, 90]µm

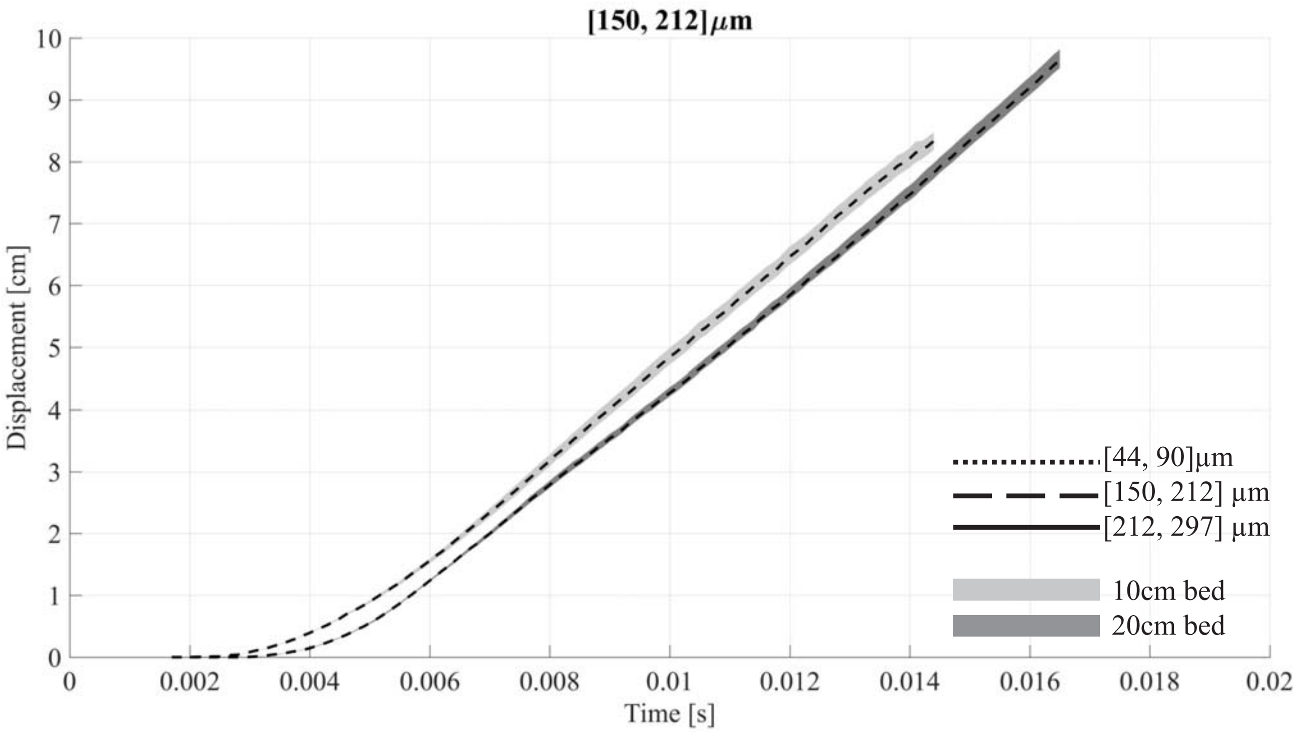
[150, 212]μm
Displacement [cm]
10
9
8
7
6
5
4
3
2
1
0
0
0.002
0.004
0.006
0.008
0.01
0.012
0.014
0.016
0.018
0.02
Time [s]
[44, 90]µm
[150, 212] µm
[212, 297] µm
10cm bed
20cm bed

CCMT

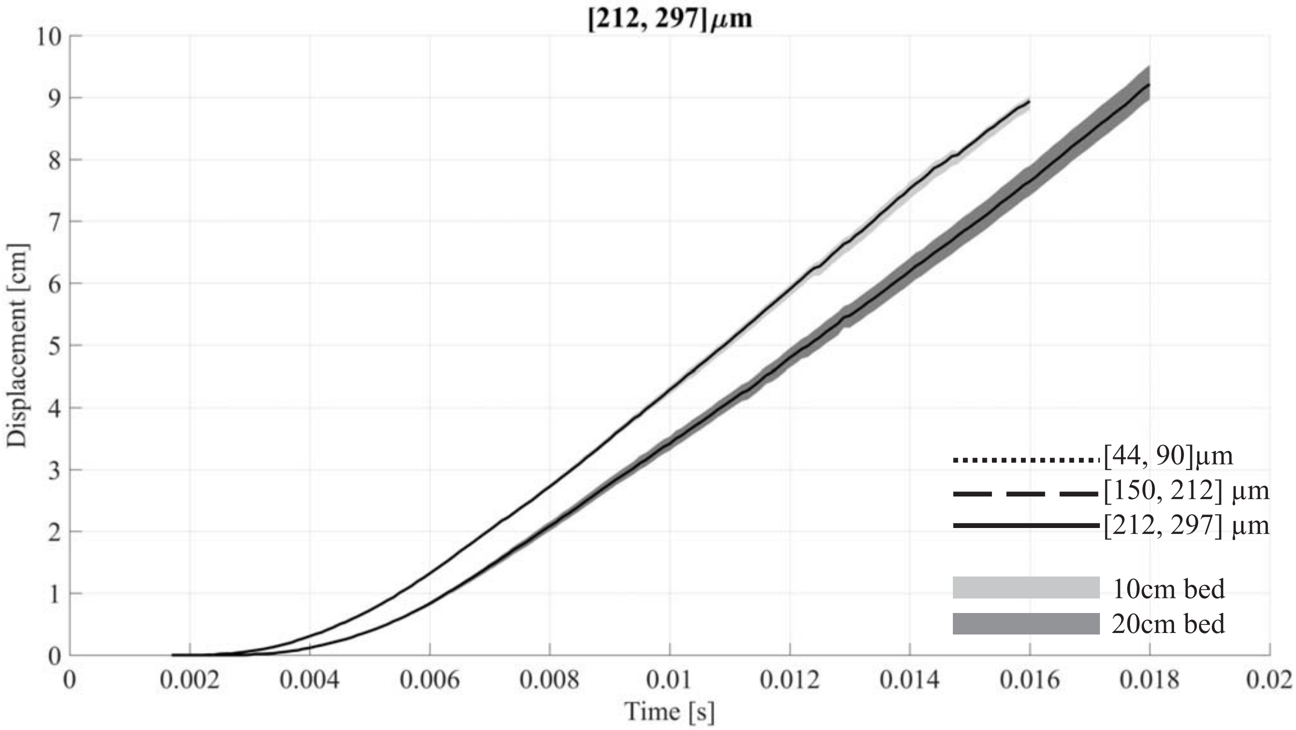
[212, 297]μm
Displacement [cm]
10
9
8
7
6
5
4
3
2
1
0
0
0.002
0.004
0.006
0.008
0.01
0.012
0.014
0.016
0.018
0.02
Time [s]
[44, 90]µm
[150, 212] µm
[212, 297] µm
10cm bed
20cm bed

CCMT

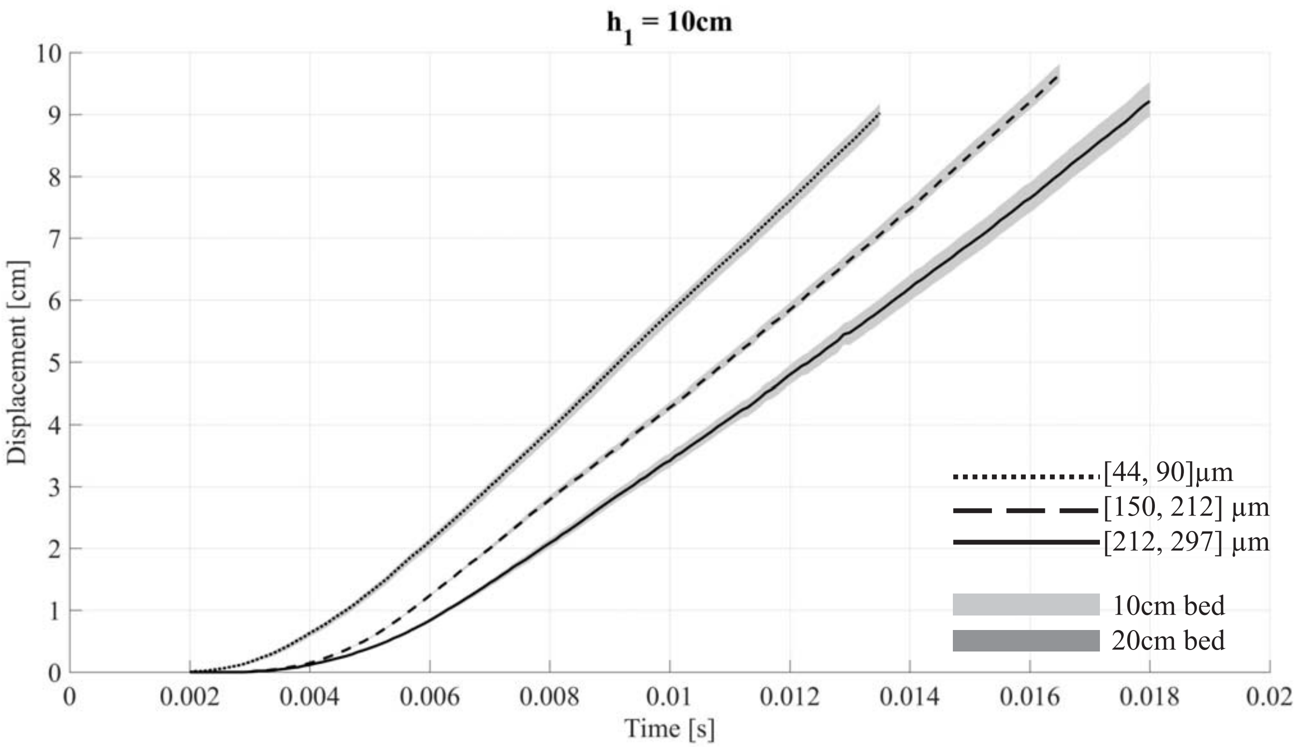

h$_1$ = 10cm

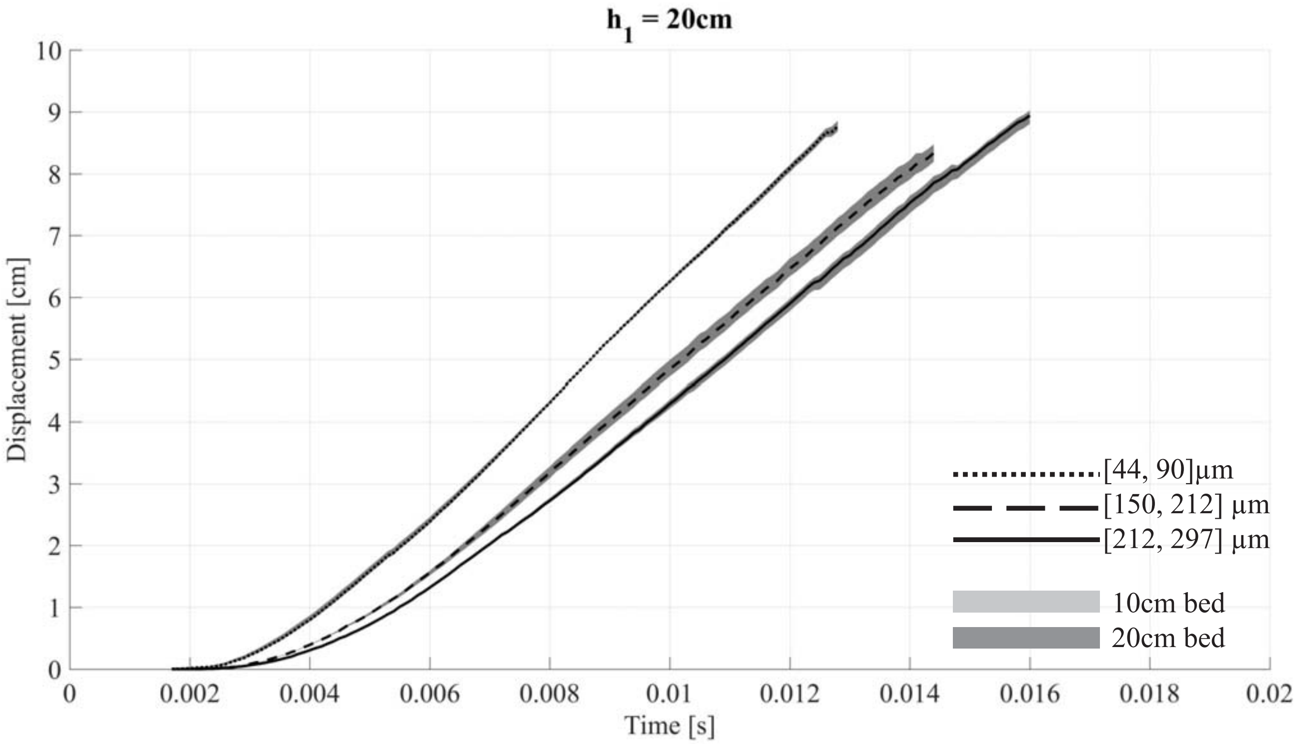

h$_1$ = 20cm

Streak Images [212, 297]µm

- Width of tube: 6 sections, each column is intensity averaged over 20 pixels

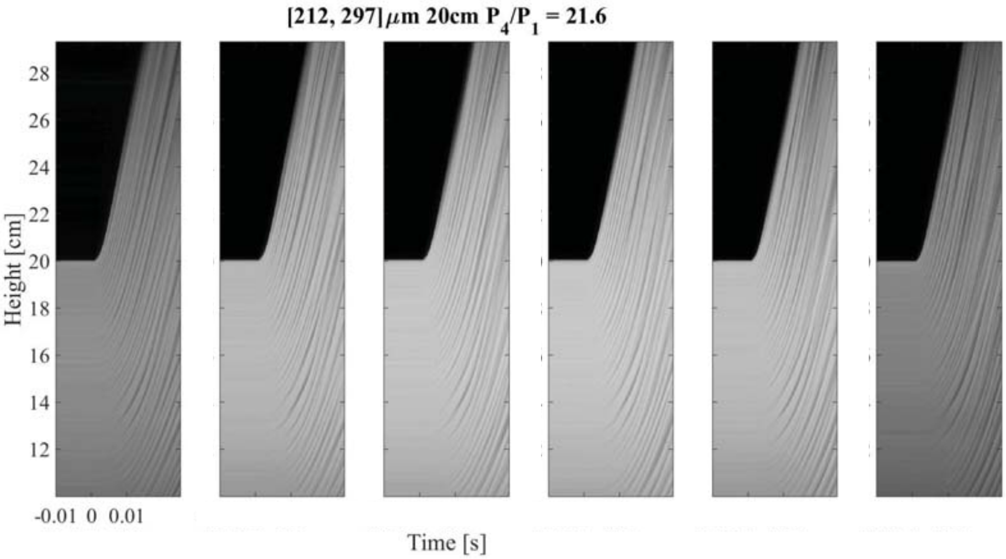

Streak Images [44, 90]µm

- Width of tube: 6 sections, each column is intensity averaged over 20 pixels

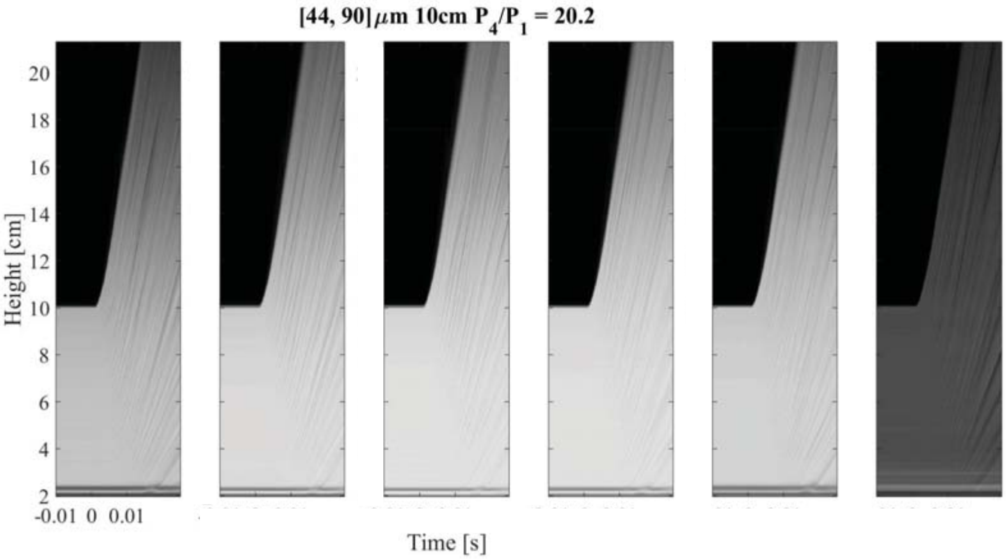

To complete the story…

- Explore various non-dimensionalizations to try to collapse the curves presented
- Develop a more complete description of the horizontal void cracks and void formations
- Compare data to available models
- Should more data be required…
  - A fourth bead diameter is available
  - A third bed height
  - For some experiments, can add in an additional pressure ratio
    - Already have some experiments with varying pressure ratios
      - Varying pressure ratio may reveal different regimes or dominating effects

CCMT

Validation Experiments: Microscale, Mesoscale, and Macroscale Experiments

Kyle Hughes, Kathy Prestridge
Los Alamos National Laboratory

Angela Diggs, Don Littrell, Mike Jenkins, Chi Mai
Munitions Directorate, Eglin AFB, FL

Nam-Ho Kim, Raphael Haftka, Chanyoung Park
University of Florida

December 13, 2018

Validation Hierarchy

- Explosive dispersal is explored at three different regimes of complexity based on the number of particles
- **Microscale** (O(1-10^2) particles)
  - Validation of microscale drag models
  - Explosive tests conducted October 2014 and February 2015
- **Mesoscale** (O(10^3 -10^6) particles)
  - Examination of instabilities, volume fraction effects, explosive dispersal at reduced scales
  - Eglin AFB explosive tests conducted November 2015
  - LANL tests conducted October 2017 and November 2018
- **Macroscale** (O(10^7)+ particles)
  - Validation of full simulation
  - AFRL blastpad conducted July 2017

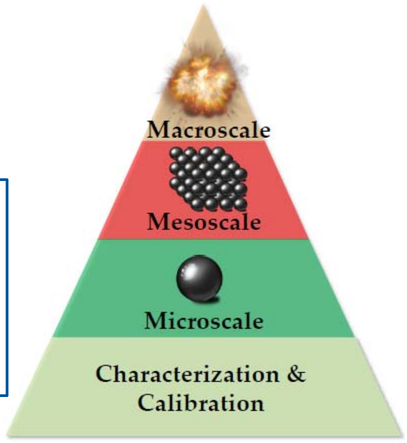

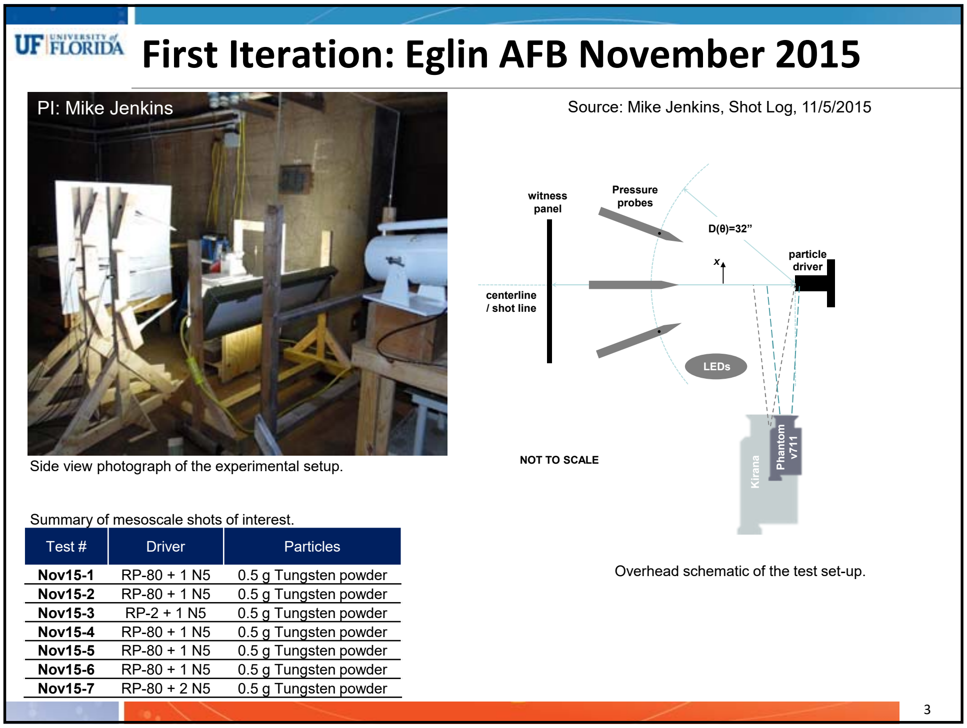

First Iteration: Eglin AFB November 2015

PI: Mike Jenkins

Side view photograph of the experimental setup.

Source: Mike Jenkins, Shot Log, 11/5/2015

Overhead schematic of the test set-up.

Summary of mesoscale shots of interest.

Test #	Driver	Particles
Nov15-1	RP-80 + 1 N5	0.5 g Tungsten powder
Nov15-2	RP-80 + 1 N5	0.5 g Tungsten powder
Nov15-3	RP-2 + 1 N5	0.5 g Tungsten powder
Nov15-4	RP-80 + 1 N5	0.5 g Tungsten powder
Nov15-5	RP-80 + 1 N5	0.5 g Tungsten powder
Nov15-6	RP-80 + 1 N5	0.5 g Tungsten powder
Nov15-7	RP-80 + 2 N5	0.5 g Tungsten powder

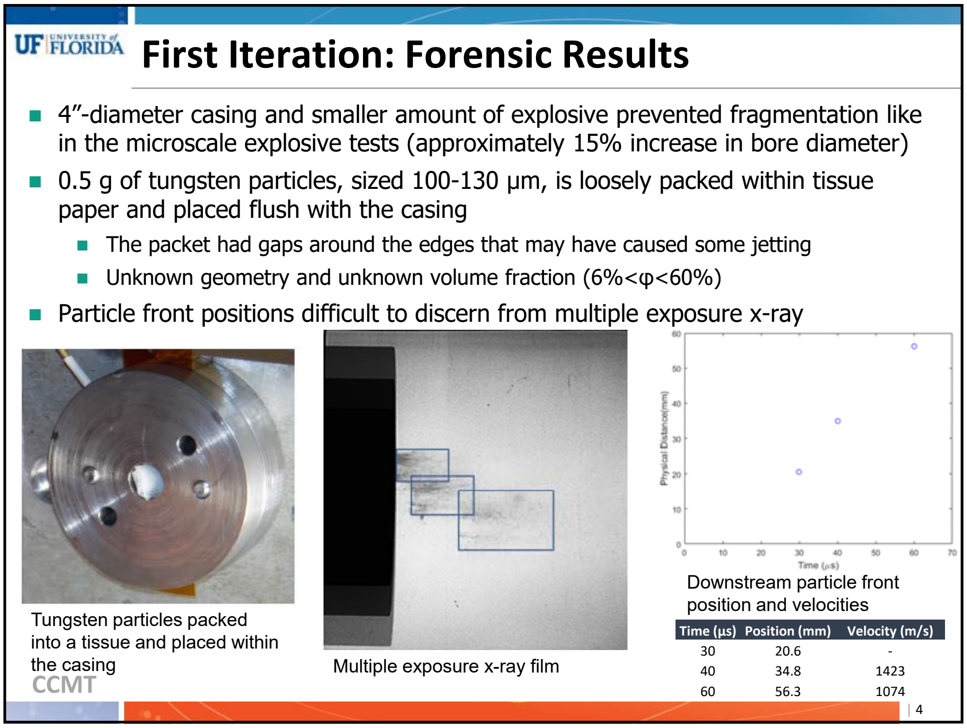

First Iteration: Forensic Results

- 4”-diameter casing and smaller amount of explosive prevented fragmentation like in the microscale explosive tests (approximately 15% increase in bore diameter)
- 0.5 g of tungsten particles, sized 100-130 μm, is loosely packed within tissue paper and placed flush with the casing
  - The packet had gaps around the edges that may have caused some jetting
  - Unknown geometry and unknown volume fraction (6%<φ<60%)
- Particle front positions difficult to discern from multiple exposure x-ray

Tungsten particles packed into a tissue and placed within the casing

Multiple exposure x-ray film

Downstream particle front position and velocities

Time (μs)	Position (mm)	Velocity (m/s)
30	20.6	-
40	34.8	1423
60	56.3	1074

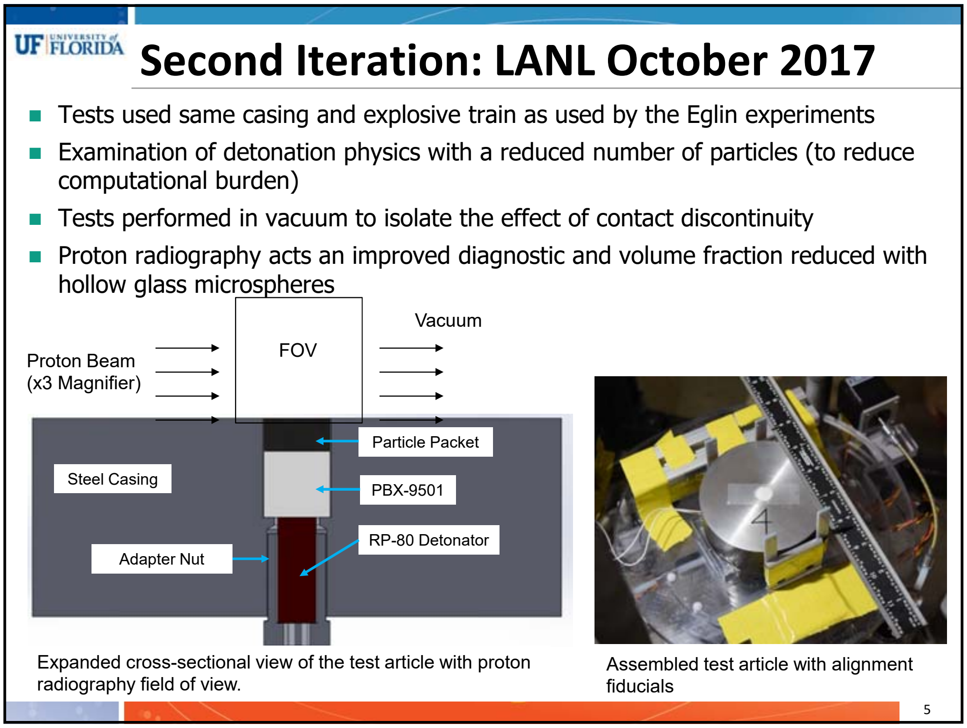

Second Iteration: LANL October 2017

- Tests used same casing and explosive train as used by the Eglin experiments
- Examination of detonation physics with a reduced number of particles (to reduce computational burden)
- Tests performed in vacuum to isolate the effect of contact discontinuity
- Proton radiography acts an improved diagnostic and volume fraction reduced with hollow glass microspheres

Expanded cross-sectional view of the test article with proton radiography field of view.

Assembled test article with alignment fiducials

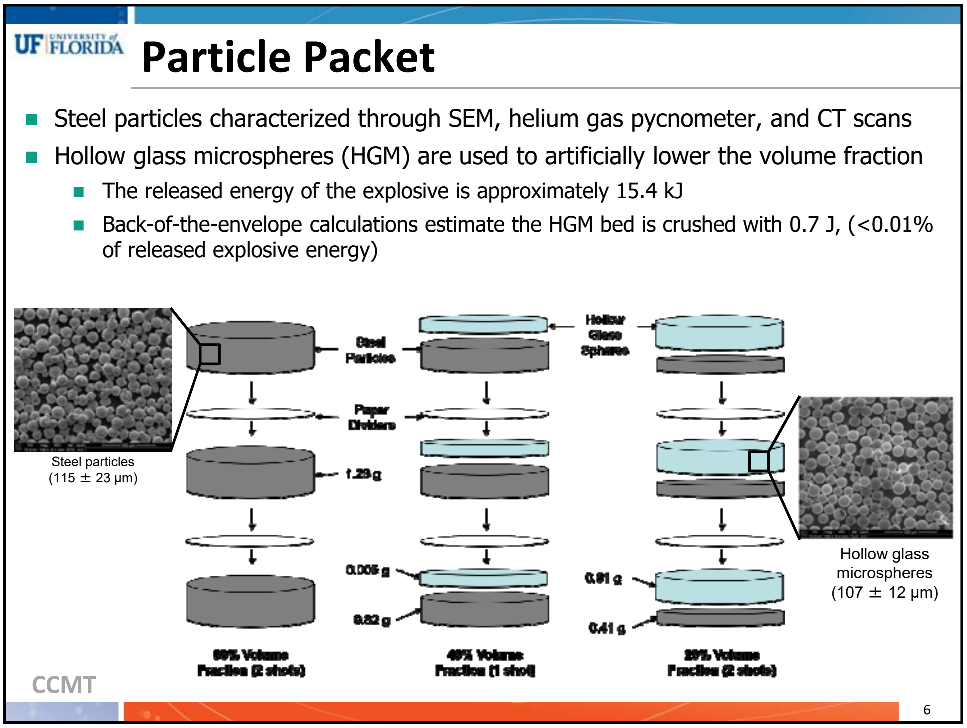

Particle Packet

- Steel particles characterized through SEM, helium gas pycnometer, and CT scans
- Hollow glass microspheres (HGM) are used to artificially lower the volume fraction
  - The released energy of the explosive is approximately 15.4 kJ
  - Back-of-the-envelope calculations estimate the HGM bed is crushed with 0.7 J, (<0.01% of released explosive energy)

Steel particles (115 ± 23 µm)

Hollow glass microspheres (107 ± 12 µm)

CCMT

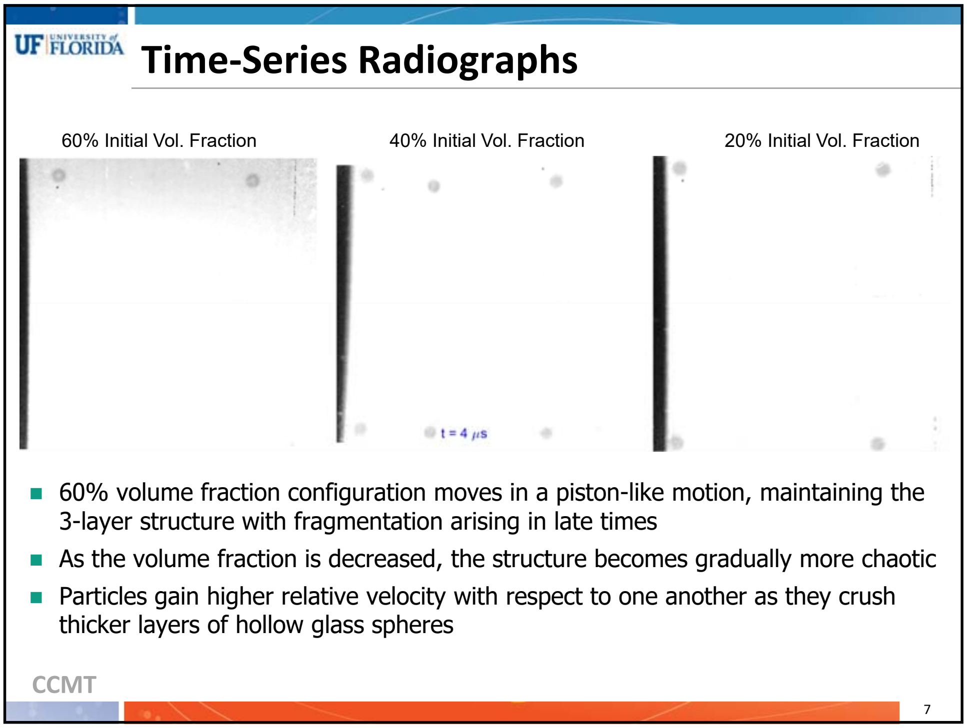

Time-Series Radiographs

- 60% volume fraction configuration moves in a piston-like motion, maintaining the 3-layer structure with fragmentation arising in late times
- As the volume fraction is decreased, the structure becomes gradually more chaotic
- Particles gain higher relative velocity with respect to one another as they crush thicker layers of hollow glass spheres

CCMT

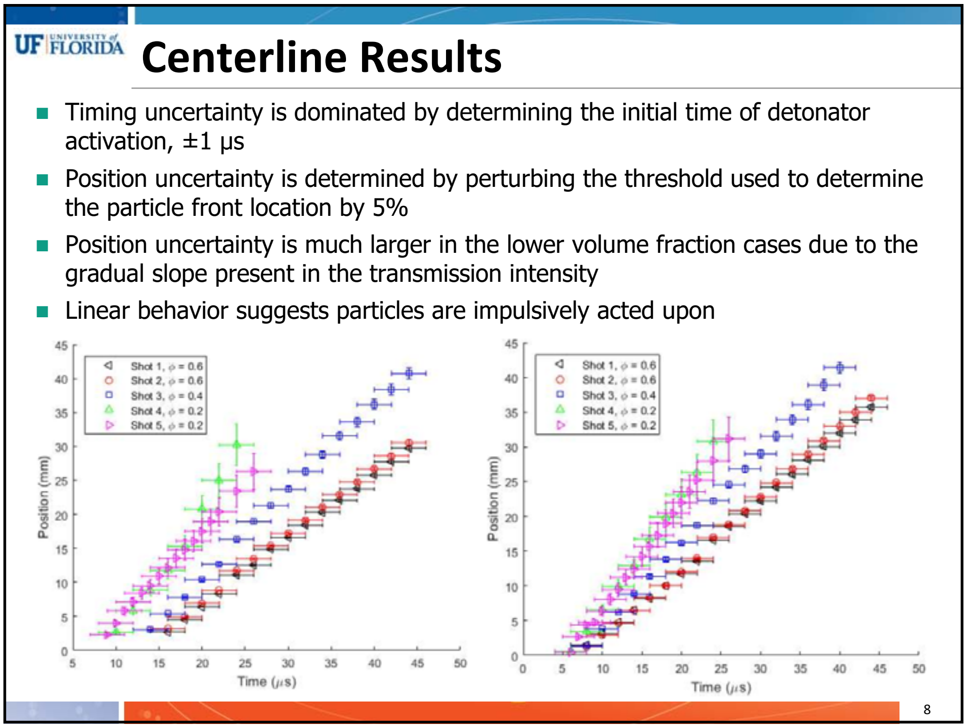

Centerline Results

- Timing uncertainty is dominated by determining the initial time of detonator activation, ±1 μs
- Position uncertainty is determined by perturbing the threshold used to determine the particle front location by 5%
- Position uncertainty is much larger in the lower volume fraction cases due to the gradual slope present in the transmission intensity
- Linear behavior suggests particles are impulsively acted upon

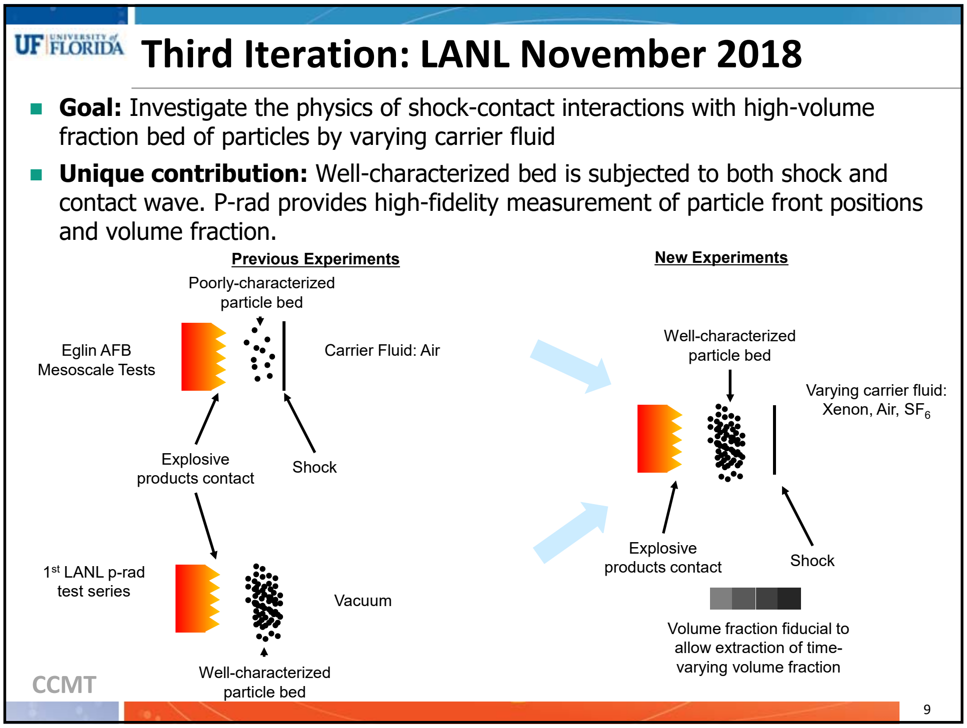

Third Iteration: LANL November 2018

- **Goal:** Investigate the physics of shock-contact interactions with high-volume fraction bed of particles by varying carrier fluid
- **Unique contribution:** Well-characterized bed is subjected to both shock and contact wave. P-rad provides high-fidelity measurement of particle front positions and volume fraction.

Previous Experiments

Poorly-characterized particle bed

Eglin AFB Mesoscale Tests

Carrier Fluid: Air

Explosive products contact

Shock

1st LANL p-rad test series

Vacuum

Well-characterized particle bed

New Experiments

Well-characterized particle bed

Varying carrier fluid: Xenon, Air, SF_6

Explosive products contact

Shock

Volume fraction fiducial to allow extraction of time-varying volume fraction

CCMT

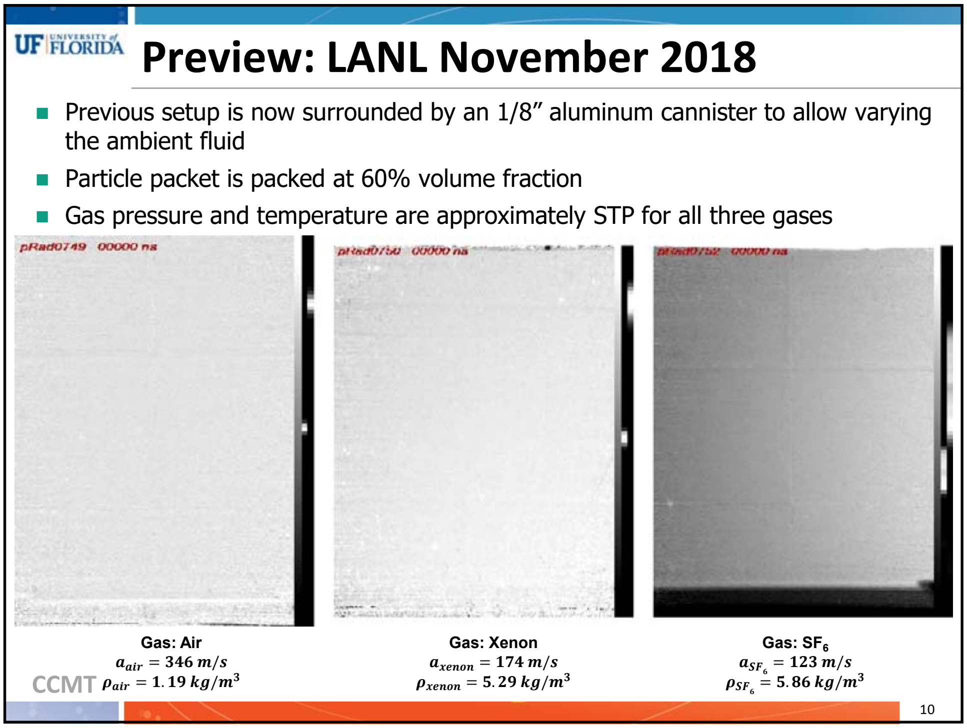

Preview: LANL November 2018

- Previous setup is now surrounded by an 1/8″ aluminum cannister to allow varying the ambient fluid
- Particle packet is packed at 60% volume fraction
- Gas pressure and temperature are approximately STP for all three gases

Gas: Air
$a_{air} = 346\ m/s$
$\rho_{air} = 1.19\ kg/m^3$

Gas: Xenon
$a_{xenon} = 174\ m/s$
$\rho_{xenon} = 5.29\ kg/m^3$

Gas: SF_6
$a_{SF_6} = 123\ m/s$
$\rho_{SF_6} = 5.86\ kg/m^3$

CCMT

Validation Hierarchy

- Explosive dispersal is explored at three different regimes of complexity based on the number of particles
- **Microscale** (O(1-10^2) particles)
  - Validation of microscale drag models
  - Explosive tests conducted October 2014 and February 2015
- **Mesoscale** (O(10^3 -10^6) particles)
  - Examination of instabilities, volume fraction effects, explosive dispersal at reduced scales
  - Eglin AFB explosive tests conducted November 2015
  - LANL tests conducted October 2017 and November 2018
- **Macroscale** (O(10^7)+ particles)
  - Validation of full simulation
  - AFRL blastpad conducted July 2017

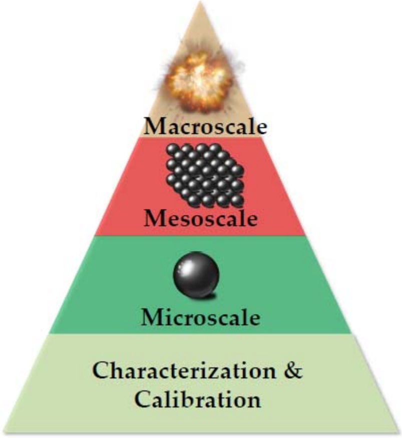

Test Facility and Planned Simulation Domain

- Eglin Blastpad chosen as the test site due to additional camera views and large number of pressure transducers (Barreto et al. 2015)
- Instrumentation suite consisting of 54 in-ground pressure transducers (sampled at 1 MHz), 6 optical linear encoders, 8 unconfined momentum traps, and 4 high-speed cameras
- Phantom v1212 sampled at 12000 fps (Camera 1/4) and Phantom v711 (Camera 2/3) sampled at 7500 fps

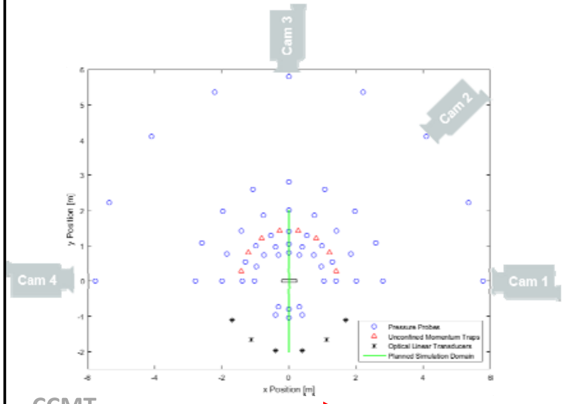

Burn direction

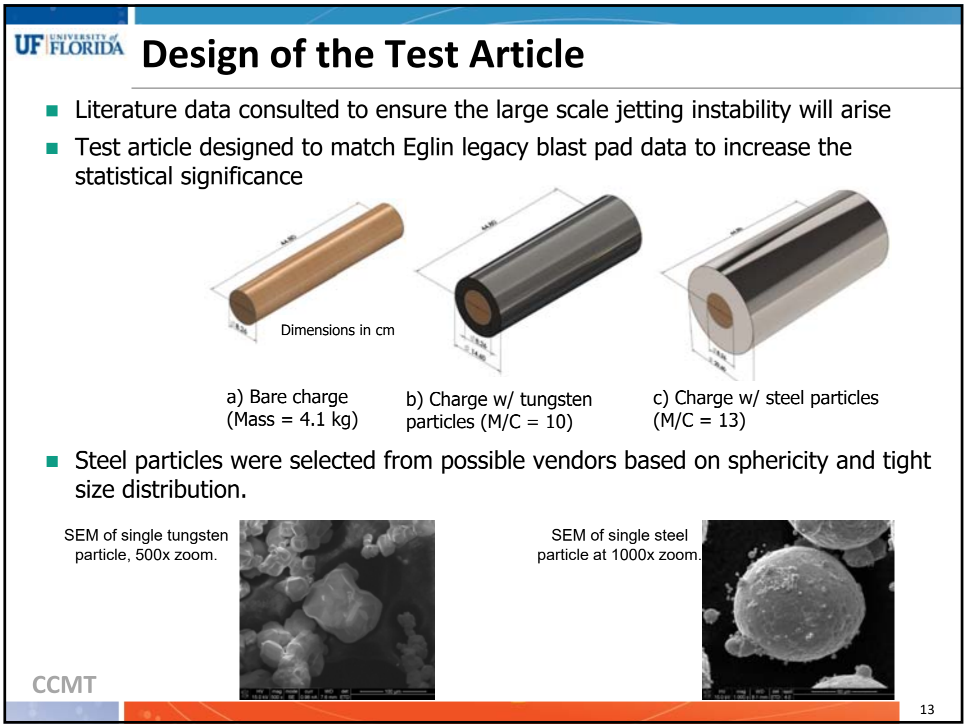

Design of the Test Article

- Literature data consulted to ensure the large scale jetting instability will arise
- Test article designed to match Eglin legacy blast pad data to increase the statistical significance

Dimensions in cm

a) Bare charge (Mass = 4.1 kg)

b) Charge w/ tungsten particles (M/C = 10)

c) Charge w/ steel particles (M/C = 13)

- Steel particles were selected from possible vendors based on sphericity and tight size distribution.

SEM of single tungsten particle, 500x zoom.

SEM of single steel particle at 1000x zoom.

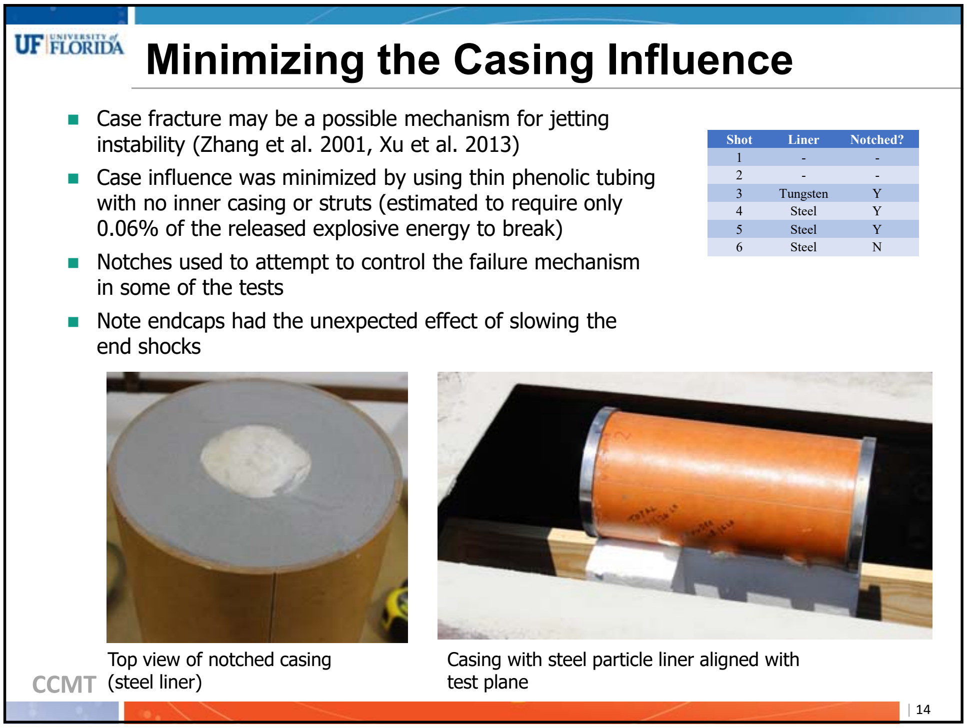

Minimizing the Casing Influence

- Case fracture may be a possible mechanism for jetting instability (Zhang et al. 2001, Xu et al. 2013)
- Case influence was minimized by using thin phenolic tubing with no inner casing or struts (estimated to require only 0.06% of the released explosive energy to break)
- Notches used to attempt to control the failure mechanism in some of the tests
- Note endcaps had the unexpected effect of slowing the end shocks

Shot	Liner	Notched?
1	-	-
2	-	-
3	Tungsten	Y
4	Steel	Y
5	Steel	Y
6	Steel	N

Top view of notched casing (steel liner)

Casing with steel particle liner aligned with test plane

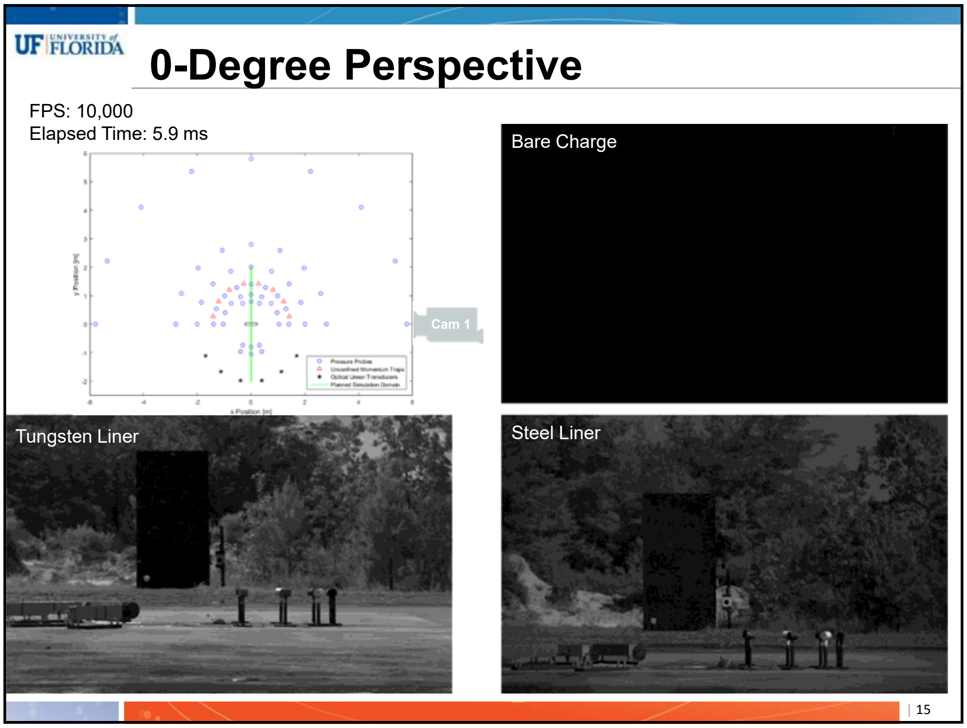

0-Degree Perspective

FPS: 10,000
Elapsed Time: 5.9 ms

Bare Charge

Tungsten Liner

Steel Liner

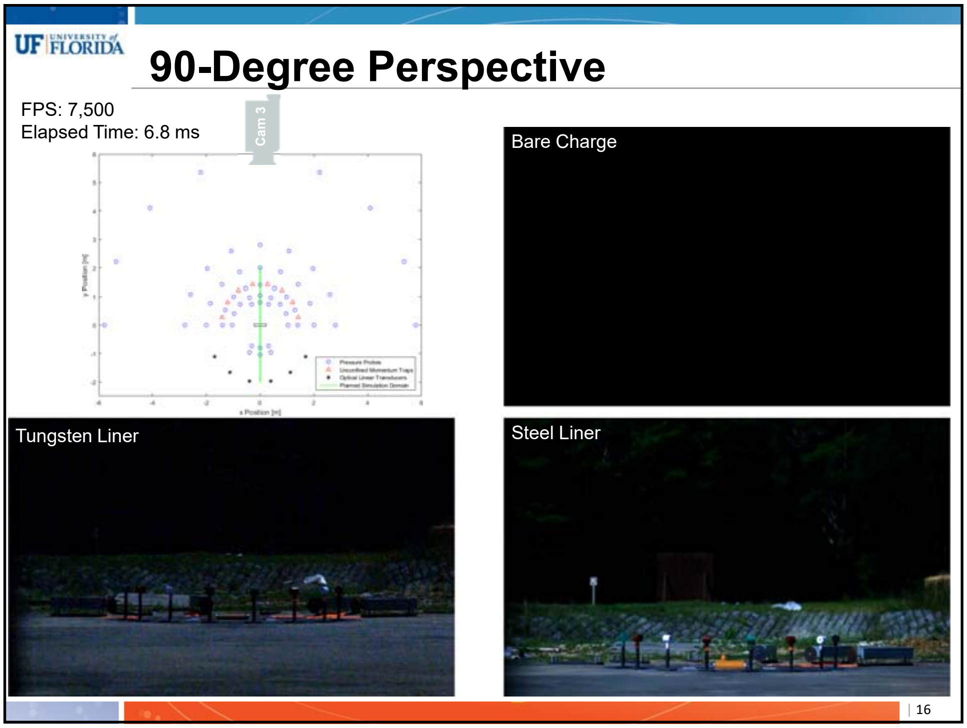

90-Degree Perspective

FPS: 7,500
Elapsed Time: 6.8 ms

Bare Charge

Tungsten Liner

Steel Liner

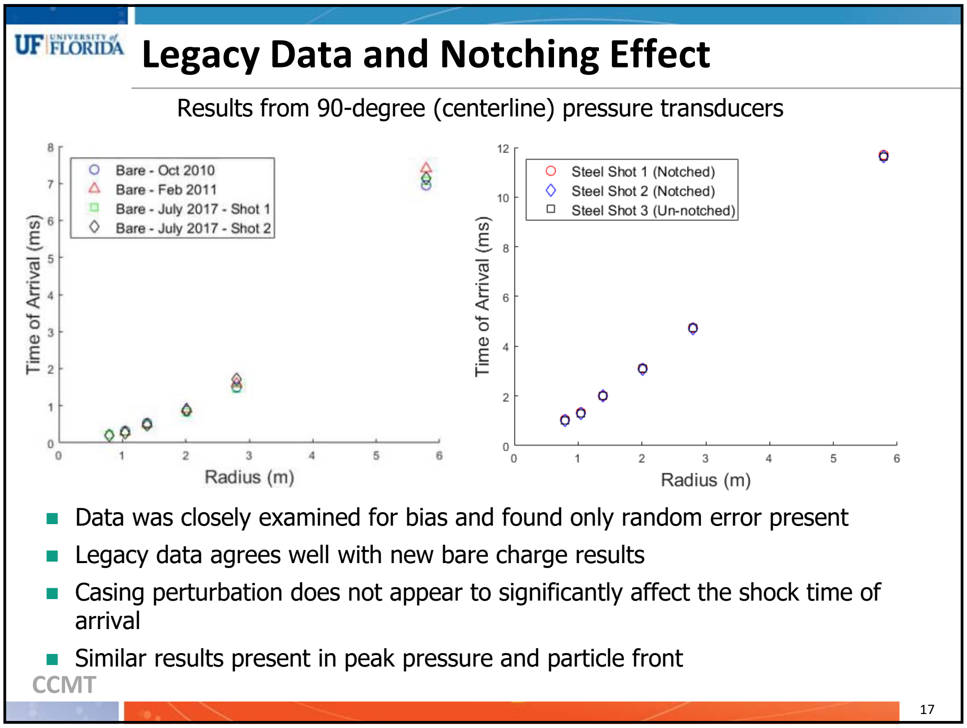

UF UNIVERSITY of FLORIDA

Legacy Data and Notching Effect

Results from 90-degree (centerline) pressure transducers

- Data was closely examined for bias and found only random error present
- Legacy data agrees well with new bare charge results
- Casing perturbation does not appear to significantly affect the shock time of arrival
- Similar results present in peak pressure and particle front

CCMT

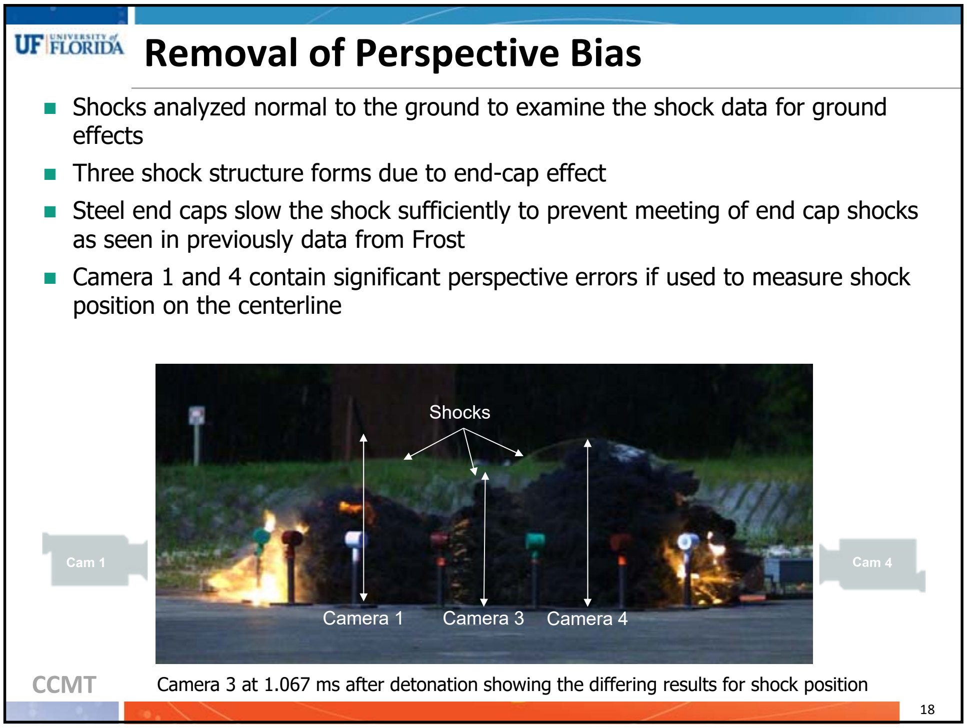

UF UNIVERSITY of FLORIDA

Removal of Perspective Bias

- Shocks analyzed normal to the ground to examine the shock data for ground effects
- Three shock structure forms due to end-cap effect
- Steel end caps slow the shock sufficiently to prevent meeting of end cap shocks as seen in previously data from Frost
- Camera 1 and 4 contain significant perspective errors if used to measure shock position on the centerline

CCMT

Camera 3 at 1.067 ms after detonation showing the differing results for shock position

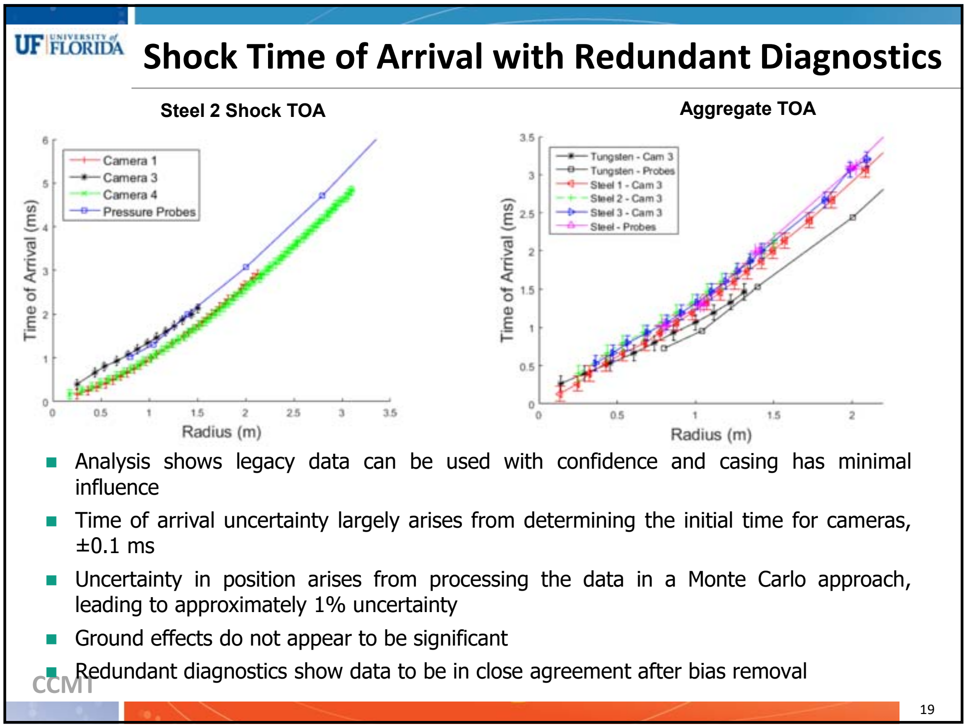

Shock Time of Arrival with Redundant Diagnostics

- Analysis shows legacy data can be used with confidence and casing has minimal influence
- Time of arrival uncertainty largely arises from determining the initial time for cameras, ±0.1 ms
- Uncertainty in position arises from processing the data in a Monte Carlo approach, leading to approximately 1% uncertainty
- Ground effects do not appear to be significant
- Redundant diagnostics show data to be in close agreement after bias removal

CCMT

19

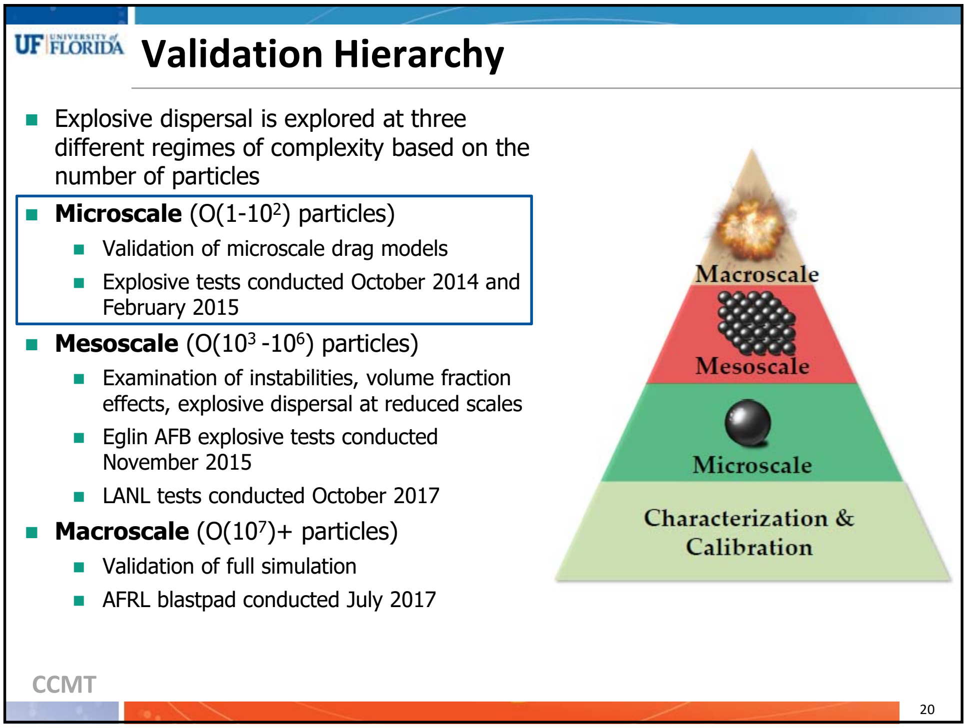

Validation Hierarchy

- Explosive dispersal is explored at three different regimes of complexity based on the number of particles
- **Microscale** ($O(1\text{-}10^2)$ particles)
  - Validation of microscale drag models
  - Explosive tests conducted October 2014 and February 2015
- **Mesoscale** ($O(10^3\text{-}10^6)$ particles)
  - Examination of instabilities, volume fraction effects, explosive dispersal at reduced scales
  - Eglin AFB explosive tests conducted November 2015
  - LANL tests conducted October 2017
- **Macroscale** ($O(10^7)$+ particles)
  - Validation of full simulation
  - AFRL blastpad conducted July 2017

CCMT

20

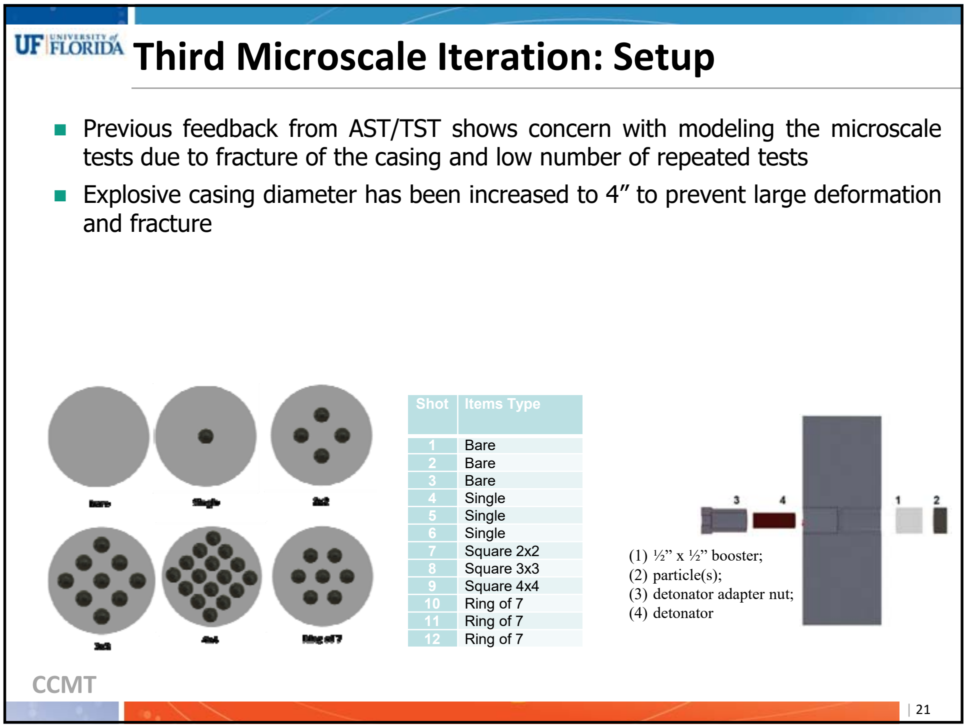

Third Microscale Iteration: Setup

- Previous feedback from AST/TST shows concern with modeling the microscale tests due to fracture of the casing and low number of repeated tests
- Explosive casing diameter has been increased to 4” to prevent large deformation and fracture

Shot	Items Type
1	Bare
2	Bare
3	Bare
4	Single
5	Single
6	Single
7	Square 2x2
8	Square 3x3
9	Square 4x4
10	Ring of 7
11	Ring of 7
12	Ring of 7

(1) ½” x ½” booster;
(2) particle(s);
(3) detonator adapter nut;
(4) detonator

CCMT

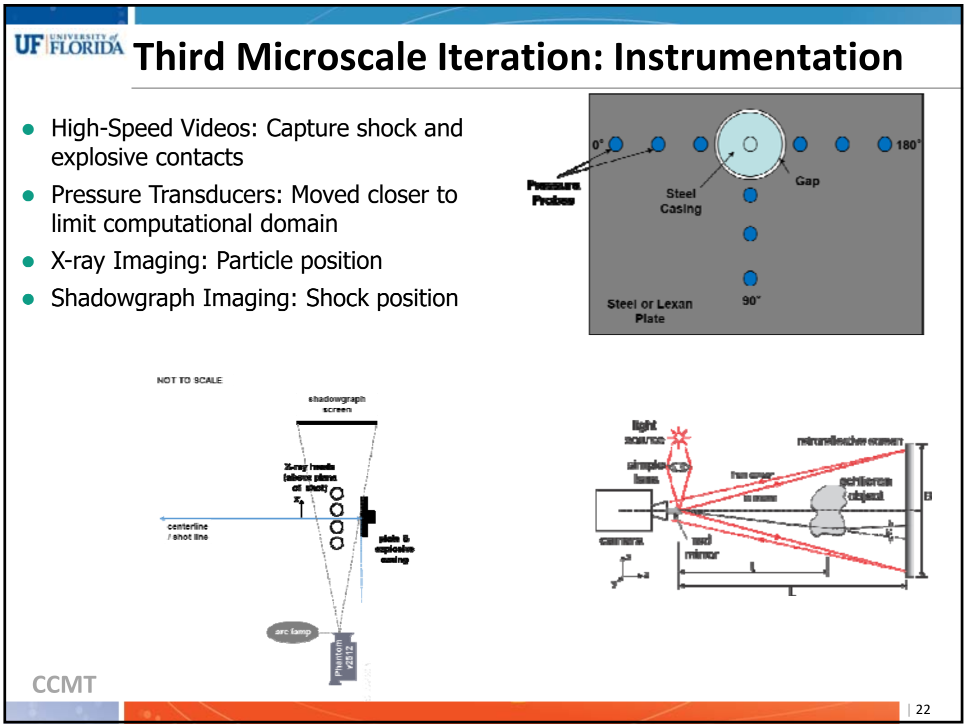

Third Microscale Iteration: Instrumentation

- High-Speed Videos: Capture shock and explosive contacts
- Pressure Transducers: Moved closer to limit computational domain
- X-ray Imaging: Particle position
- Shadowgraph Imaging: Shock position

CCMT

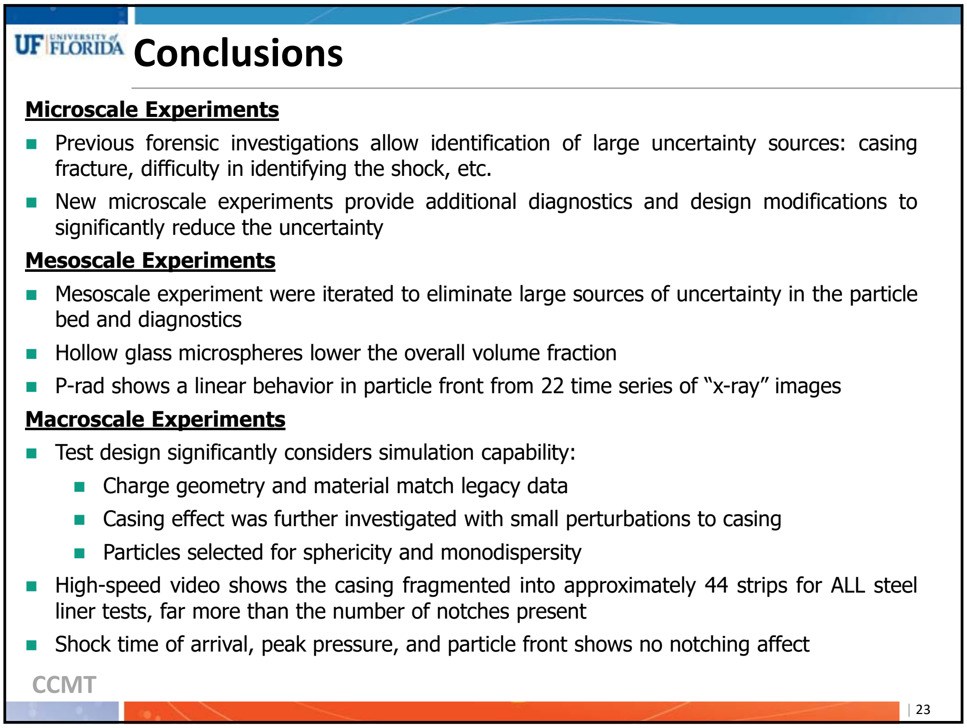

Conclusions

Microscale Experiments

- Previous forensic investigations allow identification of large uncertainty sources: casing fracture, difficulty in identifying the shock, etc.
- New microscale experiments provide additional diagnostics and design modifications to significantly reduce the uncertainty

Mesoscale Experiments

- Mesoscale experiment were iterated to eliminate large sources of uncertainty in the particle bed and diagnostics
- Hollow glass microspheres lower the overall volume fraction
- P-rad shows a linear behavior in particle front from 22 time series of "x-ray" images

Macroscale Experiments

- Test design significantly considers simulation capability:
  - Charge geometry and material match legacy data
  - Casing effect was further investigated with small perturbations to casing
  - Particles selected for sphericity and monodispersity
- High-speed video shows the casing fragmented into approximately 44 strips for ALL steel liner tests, far more than the number of notches present
- Shock time of arrival, peak pressure, and particle front shows no notching affect

CCMT

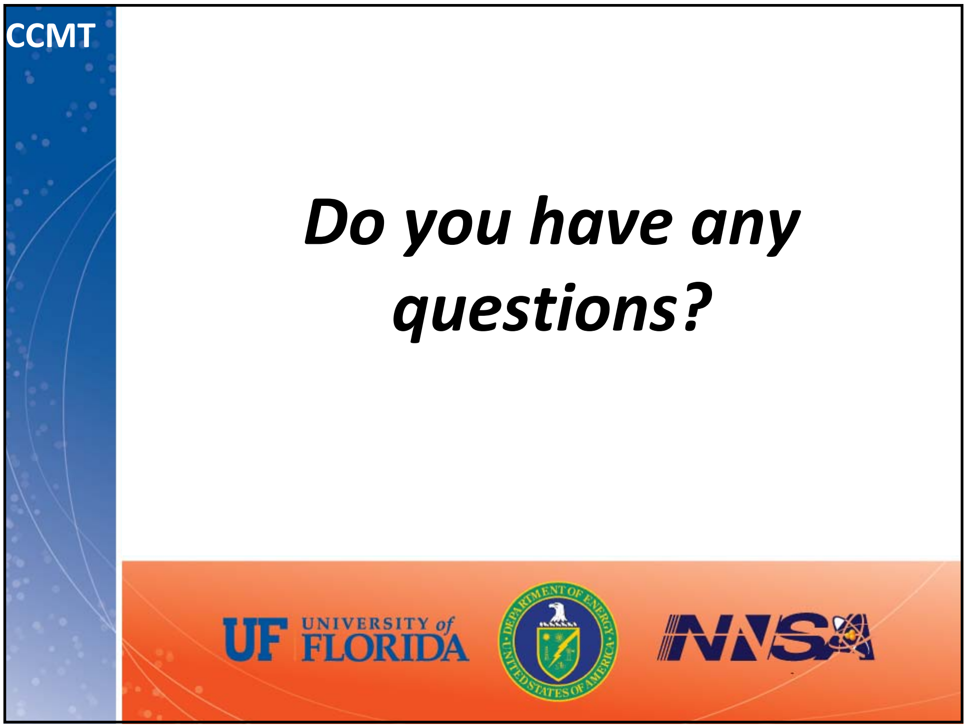

CCMT

Do you have any questions?

CCMT

CMT-nek Development and Application

Center for Compressible Multiphase Turbulence
Tri-Lab Services Team Review

December 12, 2018

Jason Hackl
David Zwick

CMT-nek users and developers

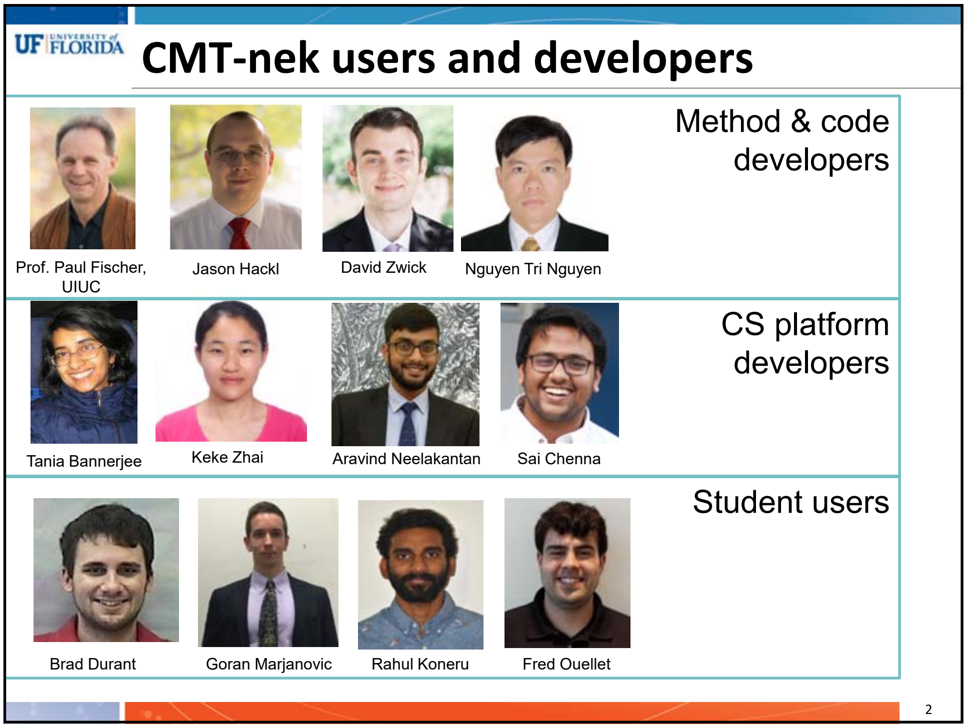

Method & code developers

Prof. Paul Fischer, UIUC
Jason Hackl
David Zwick
Nguyen Tri Nguyen

CS platform developers

Tania Bannerjee
Keke Zhai
Aravind Neelakantan
Sai Chenna

Student users

Brad Durant
Goran Marjanovic
Rahul Koneru
Fred Ouellet

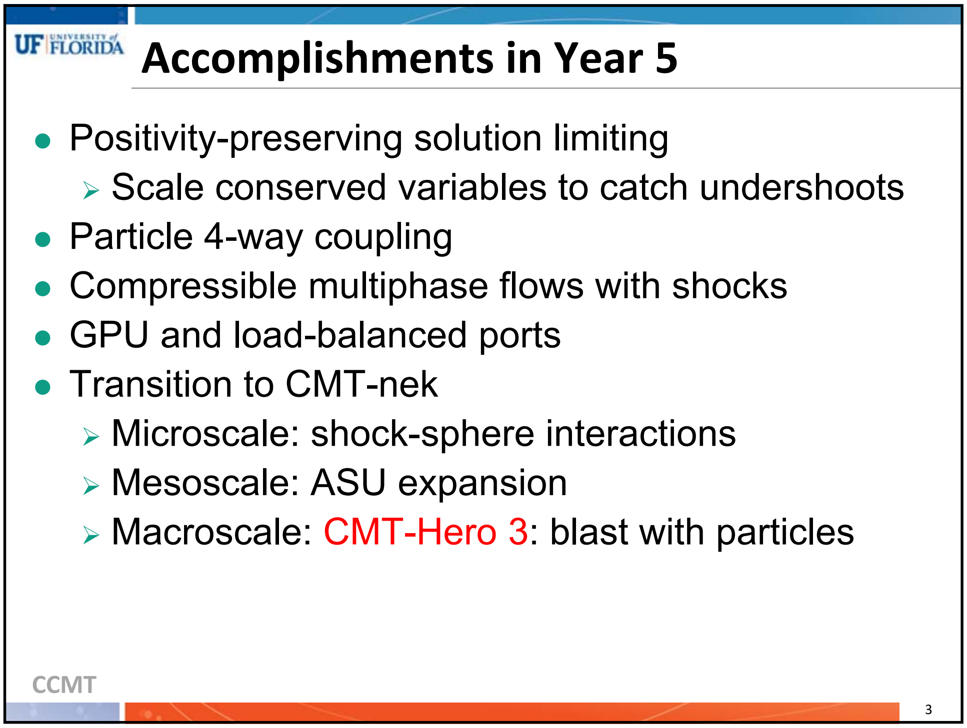

Accomplishments in Year 5

- Positivity-preserving solution limiting
  - Scale conserved variables to catch undershoots
- Particle 4-way coupling
- Compressible multiphase flows with shocks
- GPU and load-balanced ports
- Transition to CMT-nek
  - Microscale: shock-sphere interactions
  - Mesoscale: ASU expansion
  - Macroscale: CMT-Hero 3: blast with particles

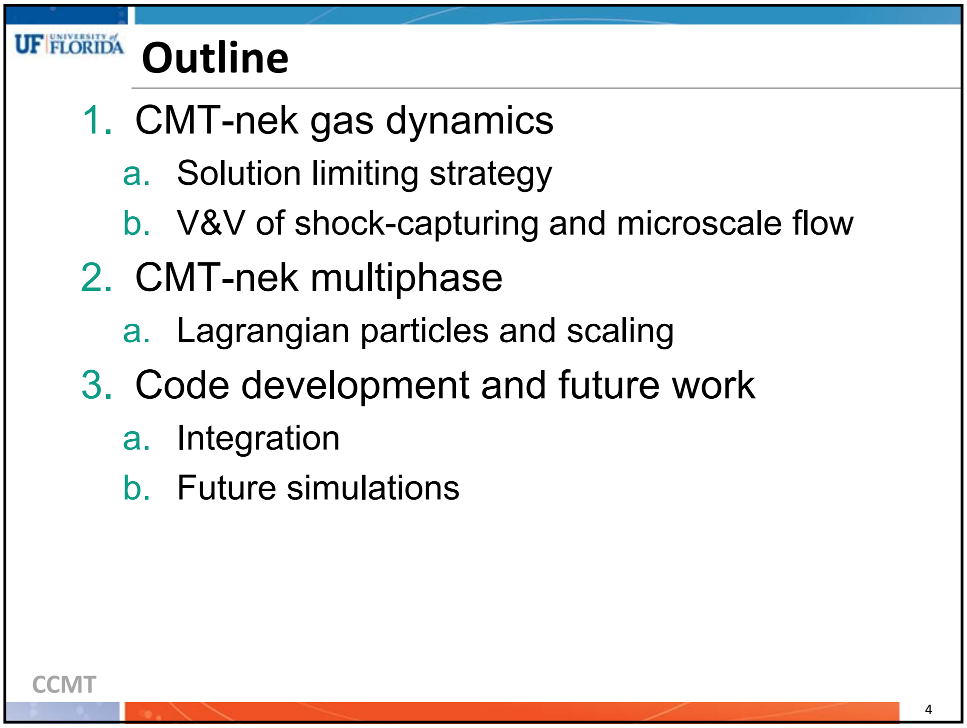

Outline

1. CMT-nek gas dynamics
   a. Solution limiting strategy
   b. V&V of shock-capturing and microscale flow
2. CMT-nek multiphase
   a. Lagrangian particles and scaling
3. Code development and future work
   a. Integration
   b. Future simulations

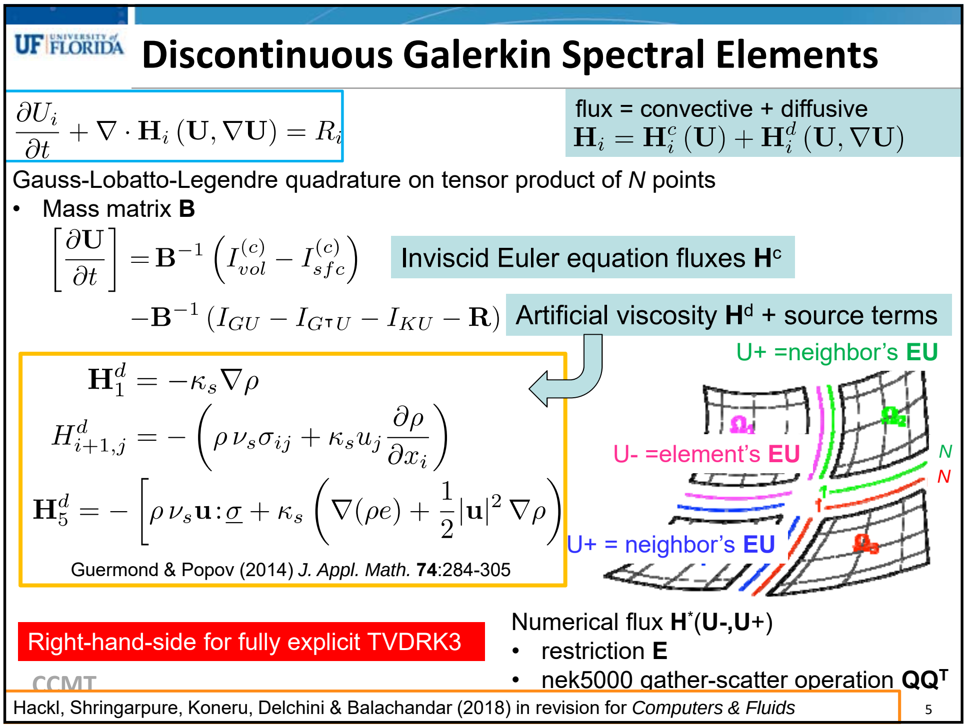

Discontinuous Galerkin Spectral Elements

$$\frac{\partial U_i}{\partial t} + \nabla \cdot \mathbf{H}_i\left(\mathbf{U}, \nabla \mathbf{U}\right) = R_i$$

flux = convective + diffusive

$$\mathbf{H}_i = \mathbf{H}_i^c\left(\mathbf{U}\right) + \mathbf{H}_i^d\left(\mathbf{U}, \nabla \mathbf{U}\right)$$

Gauss-Lobatto-Legendre quadrature on tensor product of *N* points

- Mass matrix **B**

$$\left[\frac{\partial \mathbf{U}}{\partial t}\right] = \mathbf{B}^{-1}\left(I_{vol}^{(c)} - I_{sfc}^{(c)}\right)$$

Inviscid Euler equation fluxes **H**^{c}

$$-\mathbf{B}^{-1}\left(I_{GU} - I_{G^{\intercal}U} - I_{KU} - \mathbf{R}\right)$$

Artificial viscosity **H**^{d} + source terms

$$\mathbf{H}_1^d = -\kappa_s \nabla \rho$$

$$H_{i+1,j}^d = -\left(\rho\, \nu_s \sigma_{ij} + \kappa_s u_j \frac{\partial \rho}{\partial x_i}\right)$$

$$\mathbf{H}_5^d = -\left[\rho\, \nu_s \mathbf{u} : \underline{\sigma} + \kappa_s\left(\nabla(\rho e) + \frac{1}{2}|\mathbf{u}|^2 \nabla \rho\right)\right.$$

Guermond & Popov (2014) *J. Appl. Math.* **74**:284-305

U+ =neighbor's **EU**

U- =element's **EU**

U+ = neighbor's **EU**

Right-hand-side for fully explicit TVDRK3

Numerical flux **H***(**U-**,**U+**)

- restriction **E**
- nek5000 gather-scatter operation **QQ**$^{\mathbf{T}}$

Hackl, Shringarpure, Koneru, Delchini & Balachandar (2018) in revision for *Computers & Fluids*

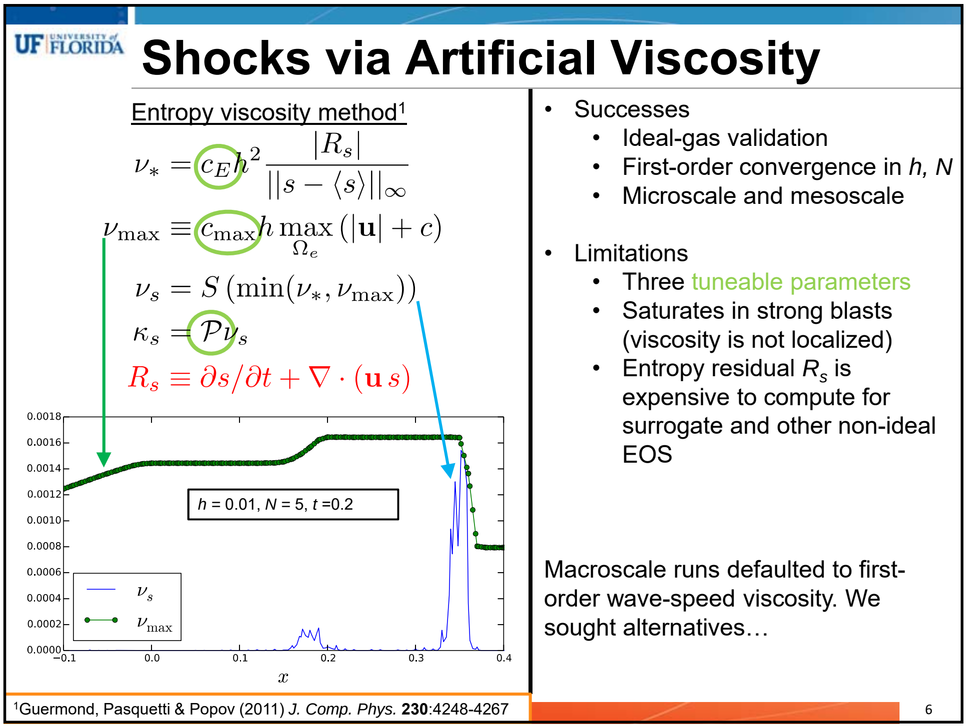

Shocks via Artificial Viscosity

Entropy viscosity method[1]

$$\nu_* = \bigcirc\!\!\!\!c_E h^2 \frac{|R_s|}{||s - \langle s \rangle||_\infty}$$

$$\nu_{\max} \equiv c_{\max} h \max_{\Omega_e}\left(|\mathbf{u}| + c\right)$$

$$\nu_s = S\left(\min(\nu_*, \nu_{\max})\right)$$

$$\kappa_s = \mathcal{P}\nu_s$$

$$R_s \equiv \partial s/\partial t + \nabla \cdot (\mathbf{u}\, s)$$

- Successes
  - Ideal-gas validation
  - First-order convergence in *h*, *N*
  - Microscale and mesoscale
- Limitations
  - Three tuneable parameters
  - Saturates in strong blasts (viscosity is not localized)
  - Entropy residual R_s is expensive to compute for surrogate and other non-ideal EOS

Macroscale runs defaulted to first-order wave-speed viscosity. We sought alternatives…

---

[1] Guermond, Pasquetti & Popov (2011) *J. Comp. Phys.* **230**:4248-4267

Internal energy limiter

[1]Zhang and Shu (2010) *JCP* **229**:8918-8934
[2]Lv & Ihme (2015) *JCP* **295**:715-739

- Rescale conserved variables ($\mathbf{U} \rightarrow^{\mathcal{L}} \mathbf{U}$) every RK stage.
  - to preserve positivity of density[1]

$$^{\mathcal{L}}\rho(r_i, s_j, t_k) = \left(\rho(r_i, s_j, t_k) - \bar{\rho}\right)\theta + \bar{\rho}$$

$$\theta = \min\left(\frac{\bar{\rho} - \epsilon}{\bar{\rho} - \rho_{\min} + \epsilon}, 1\right)$$

$$\rho_{\min} = \min_{i,j,k \in [1,N]} \left(\rho(r_i, s_j, t_k)\right)$$

  - and internal energy

$$^{\mathcal{L}}\mathbf{U} = \mathbf{U}(r_i, s_j, t_k) + \varepsilon\left(\bar{\mathbf{U}} - U(r_i, s_j, t_k)\right)$$

    - ε determined by negativity of a convex function, like internal energy, and Jensen's inequality[2].

$$\varepsilon = \frac{\tau}{\tau - \rho e\left(\bar{\mathbf{U}}\right)}, \ \tau = \min\left(0, \min_{i,j,k \in [1,N]} \rho e\left(r_i, s_j, t_k\right)\right)$$

$$\bar{\rho} = \int_{\Omega_e} \rho dV \approx \sum_{i,j,k=1}^{N} \omega_i \omega_j \omega_k \rho\left(r_i, s_j, t_k\right)$$

Shock-sphere drag coefficient

- Mach 3 shock-sphere interaction, N = 5

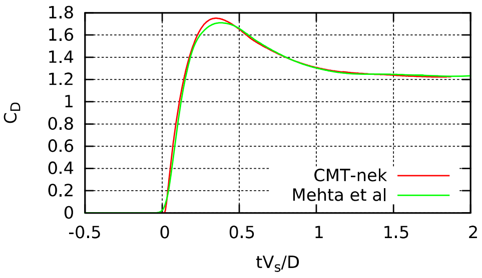

- Mach 1.22 shock-sphere, two resolutions

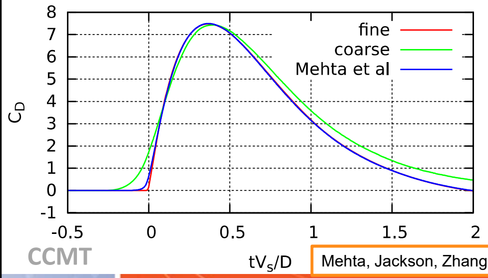

	Mach 1.22	Mach 3
$c_{\max}$	0.3	0.75
c_E	40	85
P	0.75	0.75

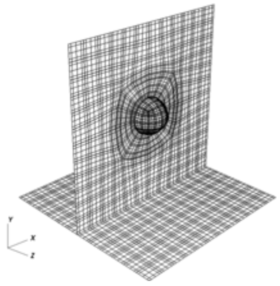

Coarse: 2048 elements, N = 5
Fine: 1 M elements, N = 5

CCMT

Mehta, Jackson, Zhang & Balachandar (2016) *J. Appl. Phys.* **119**:104901

Shock-sphere interaction

Mach 1.22

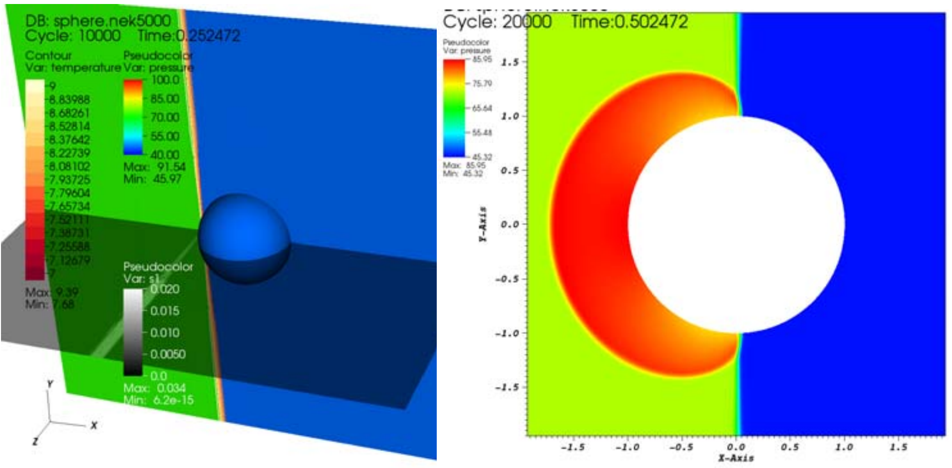

Captured features identified in experiments

- Tanno *et al* (2003) *Shock Waves* **13**(3):191-200

CCMT

Microscale: Rarefaction over spheres

Tests generalized Faxen theorem[1]

$$\mathbf{F}^{\mathrm{GF}}(t) = \frac{4}{3}\pi R^3 \rho^{\mathrm{un}} \overline{\frac{\mathrm{D}\mathbf{u}}{\mathrm{D}t}^{\mathrm{un}}}^{V} + 4\pi R^3 \int_{\tilde{\xi}=-\infty}^{\tilde{t}} K_{\mathrm{iu}}(\tilde{t}-\tilde{\xi}; M)\left[\frac{\mathrm{D}}{\mathrm{D}t}\overline{(\rho\mathbf{u}_r)^{\mathrm{un}}}^{S}\right]_{\tilde{t}=\tilde{\xi}} \mathrm{d}\tilde{\xi}$$

N_{elem}	30 K
DOF	65 M
Procs	8192

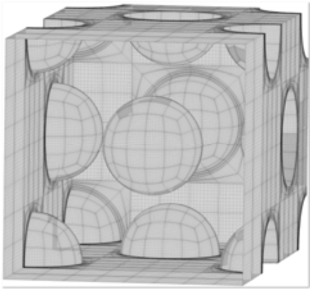

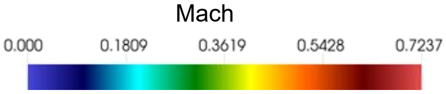

FCC 10% Volume fraction

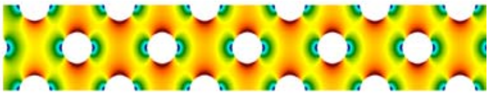

FCC 15% Volume fraction

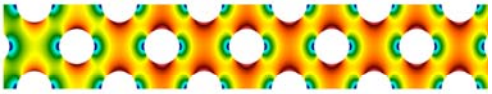

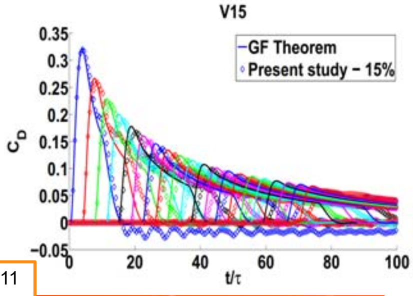

[1]Annamalai & Balachandar (2017) *J. Fluid. Mech.* **816**, 381-411
Marjanovic *et al* (2018) *Phys. Rev. Fluids* **3**, 094301

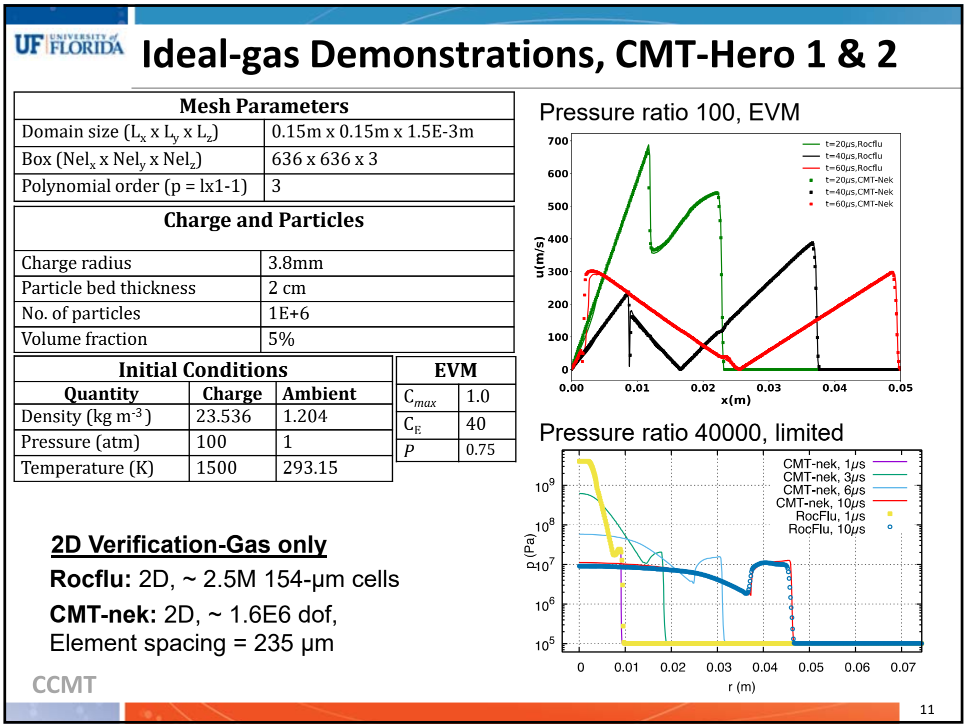

Ideal-gas Demonstrations, CMT-Hero 1 & 2

Mesh Parameters	
Domain size (L_x x L_y x L_z)	0.15m x 0.15m x 1.5E-3m
Box (Nel_x x Nel_y x Nel_z)	636 x 636 x 3
Polynomial order (p = lx1-1)	3

Charge and Particles	
Charge radius	3.8mm
Particle bed thickness	2 cm
No. of particles	1E+6
Volume fraction	5%

Initial Conditions		
Quantity	Charge	Ambient
Density (kg m^{-3})	23.536	1.204
Pressure (atm)	100	1
Temperature (K)	1500	293.15

EVM	
C_{max}	1.0
C_E	40
P	0.75

2D Verification-Gas only

Rocflu: 2D, ~ 2.5M 154-μm cells

CMT-nek: 2D, ~ 1.6E6 dof, Element spacing = 235 μm

Pressure ratio 100, EVM

Pressure ratio 40000, limited

CCMT

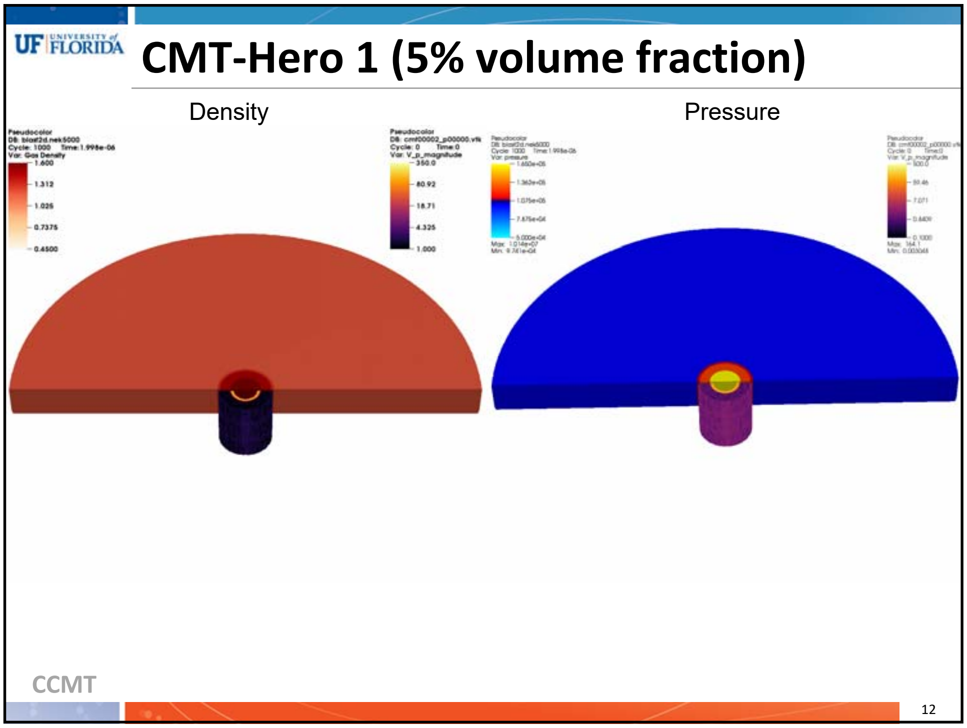

CMT-Hero 1 (5% volume fraction)

Density

Pressure

CCMT

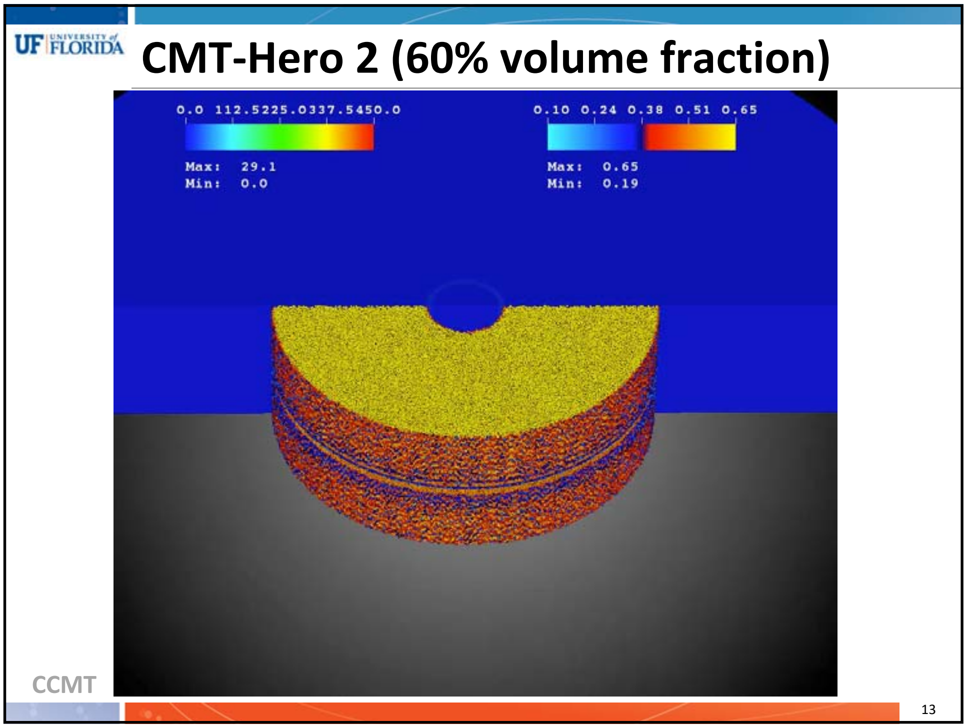

CMT-Hero 2 (60% volume fraction)

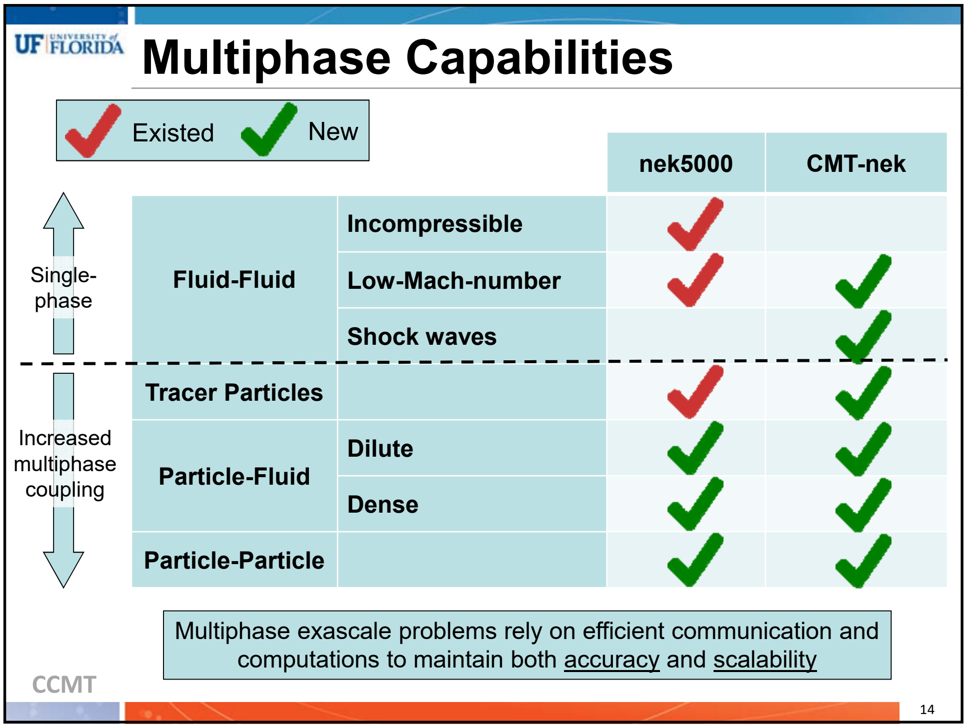

Multiphase Capabilities

✔ Existed ✓ New

		nek5000	CMT-nek
Fluid-Fluid	**Incompressible**	✔	
	Low-Mach-number	✔	✓
	Shock waves		✓
Tracer Particles		✔	✓
Particle-Fluid	**Dilute**	✓	✓
	Dense	✓	✓
Particle-Particle		✓	✓

Multiphase exascale problems rely on efficient communication and computations to maintain both accuracy and scalability

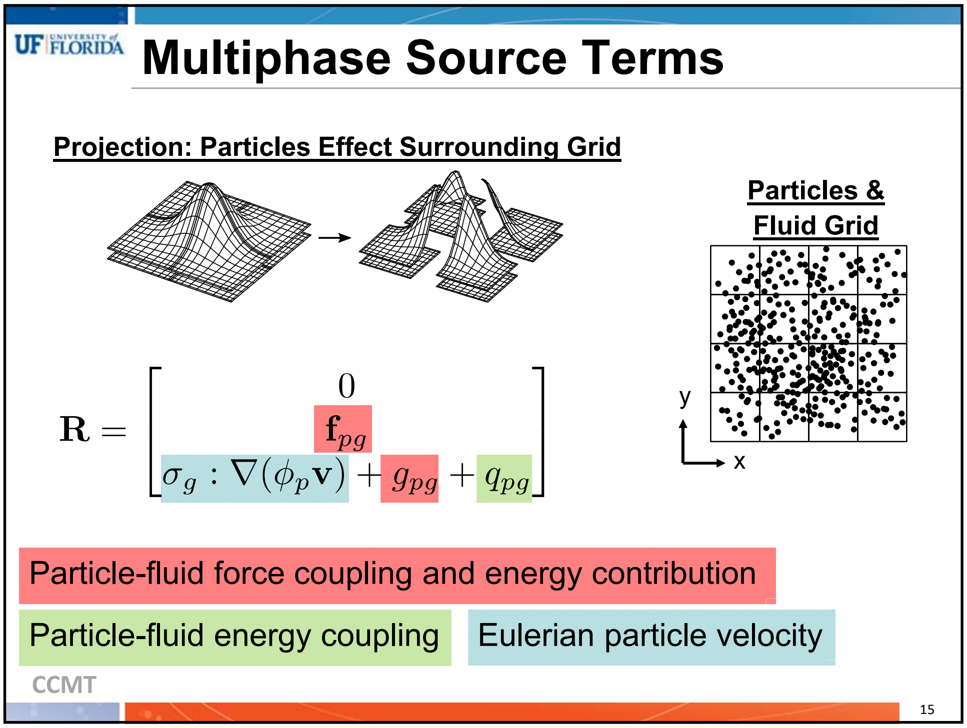

Multiphase Source Terms

Projection: Particles Effect Surrounding Grid

Particles & Fluid Grid

$$\mathbf{R} = \begin{bmatrix} 0 \\ \mathbf{f}_{pg} \\ \sigma_g : \nabla(\phi_p \mathbf{v}) + g_{pg} + q_{pg} \end{bmatrix}$$

🟥 Particle-fluid force coupling and energy contribution

🟩 Particle-fluid energy coupling

🟦 Eulerian particle velocity

CCMT

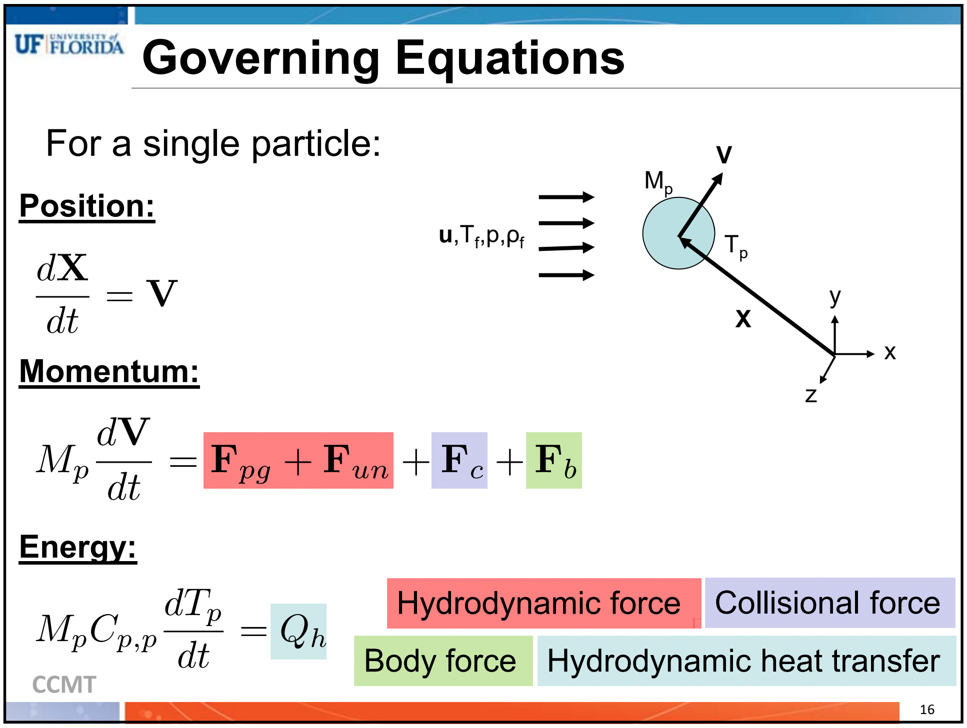

Governing Equations

For a single particle:

Position:

$$\frac{d\mathbf{X}}{dt} = \mathbf{V}$$

Momentum:

$$M_p \frac{d\mathbf{V}}{dt} = \mathbf{F}_{pg} + \mathbf{F}_{un} + \mathbf{F}_c + \mathbf{F}_b$$

Energy:

$$M_p C_{p,p} \frac{dT_p}{dt} = Q_h$$

🟥 Hydrodynamic force

🟪 Collisional force

🟩 Body force

🟦 Hydrodynamic heat transfer

CCMT

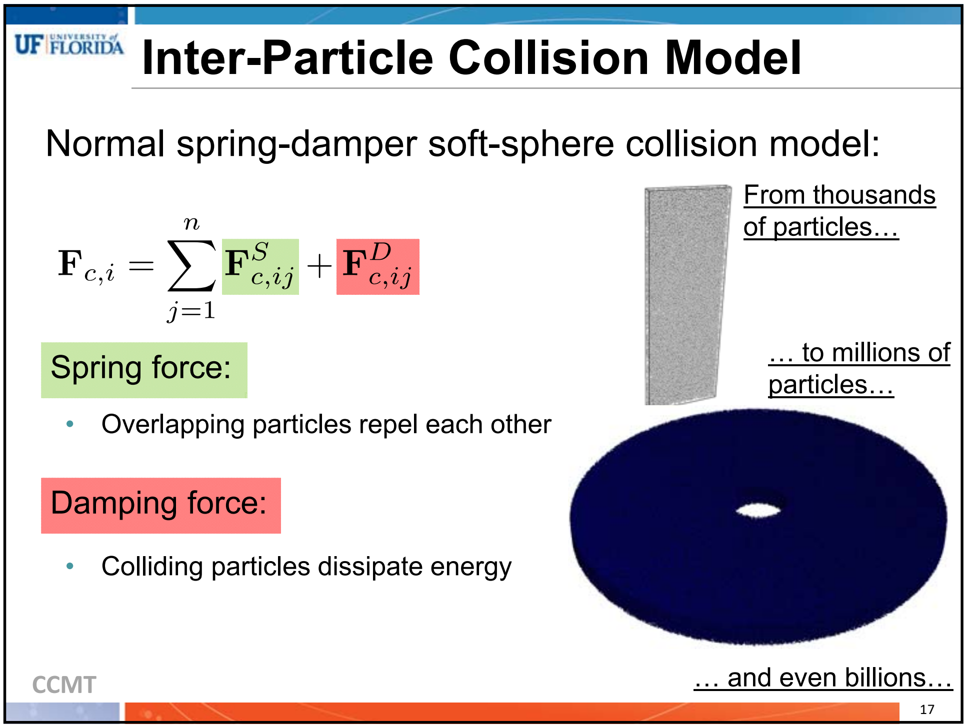

Inter-Particle Collision Model

Normal spring-damper soft-sphere collision model:

$$\mathbf{F}_{c,i} = \sum_{j=1}^{n} \mathbf{F}^{S}_{c,ij} + \mathbf{F}^{D}_{c,ij}$$

Spring force:

- Overlapping particles repel each other

Damping force:

- Colliding particles dissipate energy

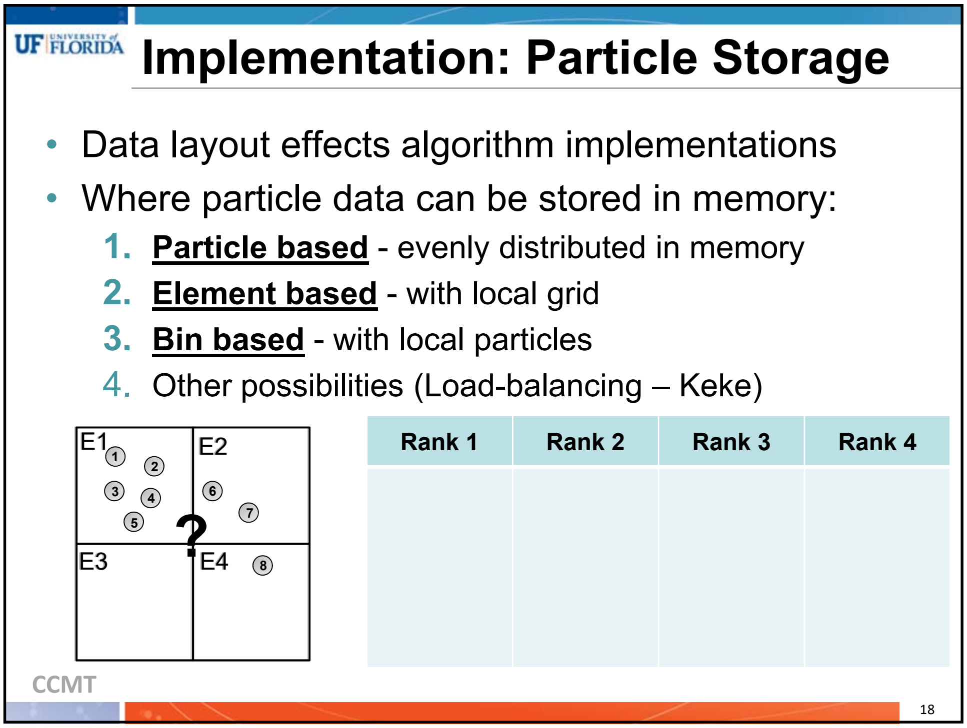

Implementation: Particle Storage

- Data layout effects algorithm implementations
- Where particle data can be stored in memory:
  1. **Particle based** - evenly distributed in memory
  2. **Element based** - with local grid
  3. **Bin based** - with local particles
  4. Other possibilities (Load-balancing – Keke)

Rank 1	Rank 2	Rank 3	Rank 4

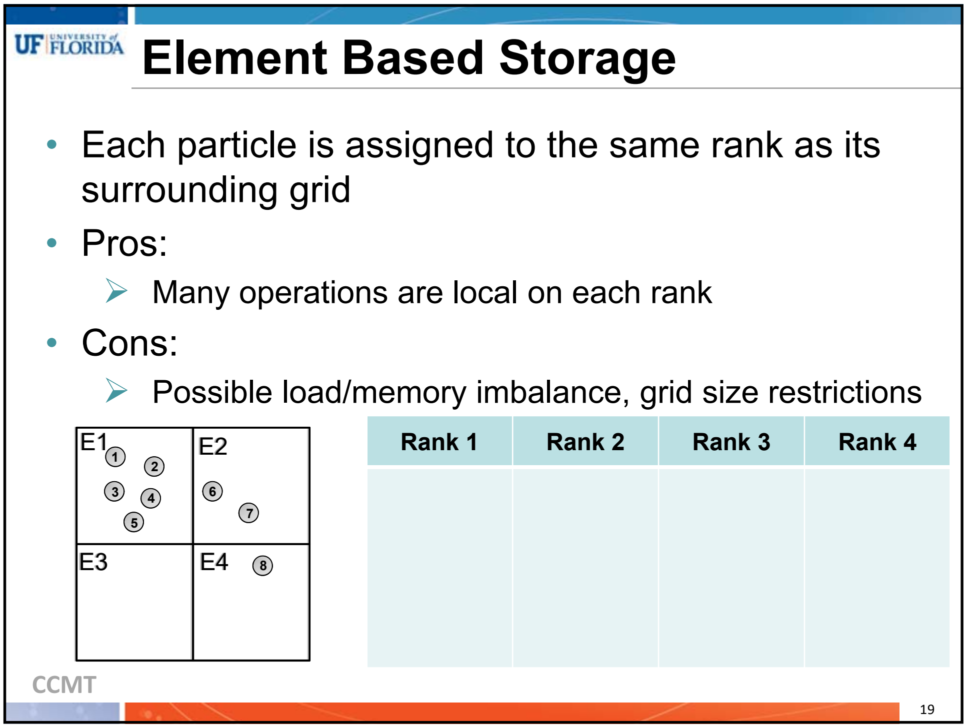

Element Based Storage

- Each particle is assigned to the same rank as its surrounding grid
- Pros:
  - Many operations are local on each rank
- Cons:
  - Possible load/memory imbalance, grid size restrictions

Rank 1	Rank 2	Rank 3	Rank 4

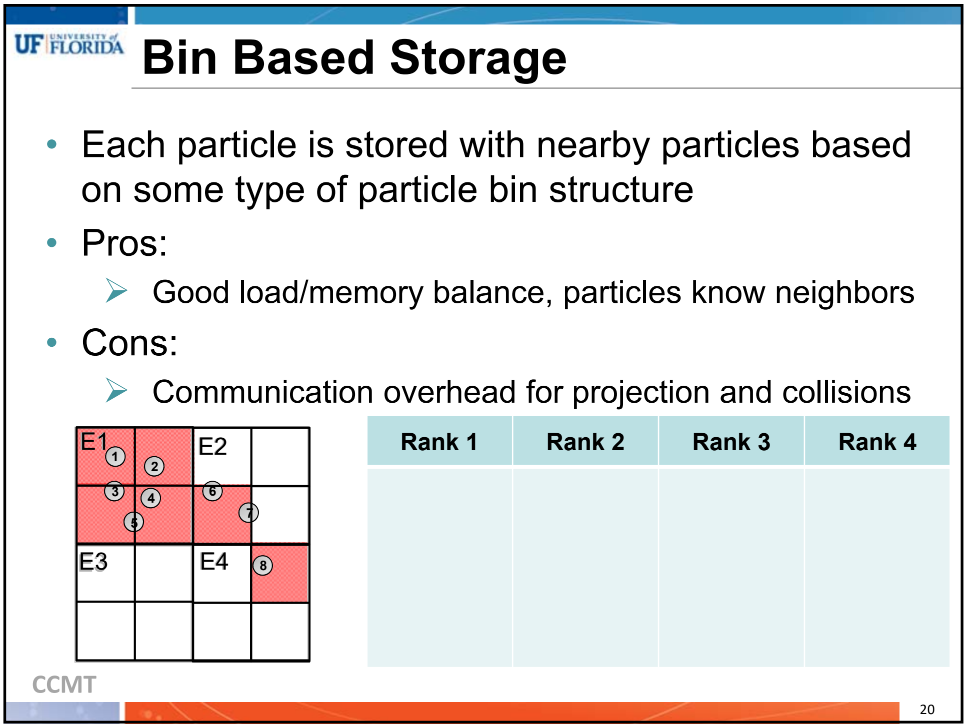

Bin Based Storage

- Each particle is stored with nearby particles based on some type of particle bin structure
- Pros:
  - Good load/memory balance, particles know neighbors
- Cons:
  - Communication overhead for projection and collisions

Rank 1	Rank 2	Rank 3	Rank 4

Scaling

How does bin based algorithm scale?

Setup:

- Uniform flow in a 3D box
- Particles are randomly distributed in box
- Cases run on Vulcan (LLNL)
- Bin based is better than element based storage
- Scaling:
  - Parameter scaling (N_p, N, δ_f)
  - Processor scaling (strong and weak scaling)

Bin Based Processor Scaling

Strong Scaling

ASU SimulationDAT Weak Scaling

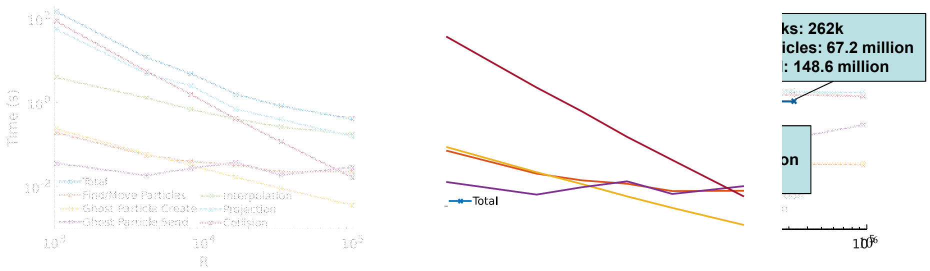

Description:

- Problem fixed while R processors varied
- 25 million total particles
- N_p/R decreases from 24,500 to 256

- Problem per rank fixed while R processors varied
- 3,072 particles per rank
- N_p increases from 3 to 302 million

Observations:

- Collision algorithm scales as $O((N_p/R)^2)$
- Creating ghost particles scales as $O(N_p/R)$

Observations:

- Sending ghost particles scales as $O(\log_2 R)$

Fluidization

- Particle bed initially at ~ ¼ of height
- Particles initially randomly distributed in entire domain and allowed to settle

Parameter	Value
N_p	9,240
D_p [m]	1,200E-6
ρ_p [kg/m^3]	1,000
L_x [m]	0.044
L_y [m]	0.120
L_z [m]	0.010
ρ_f [kg/m^3]	1.205
v_f [Pa s]	1.8E-5

Run	Inlet Superficial Velocity [m/s]
1	0.083
2	0.165
3	0.248
4	0.331
5	0.414
6	0.496

Experimental minimum fluidization velocity ~ 0.3

Mesh and Boundary Conditions

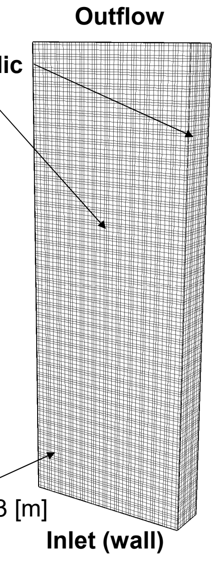

Pressure Drop

- Superficial velocity is intrinsic fluid velocity multiplied by fluid volume fraction at inlet
- Average pressure drop across bed is normalized by bed weight

Average pressure drop vs. inlet velocity

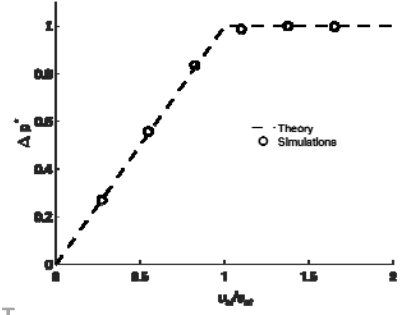

Particle volume fraction (Run 6)

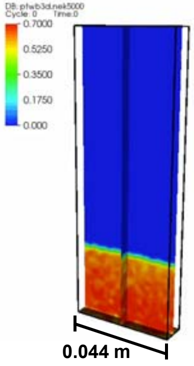

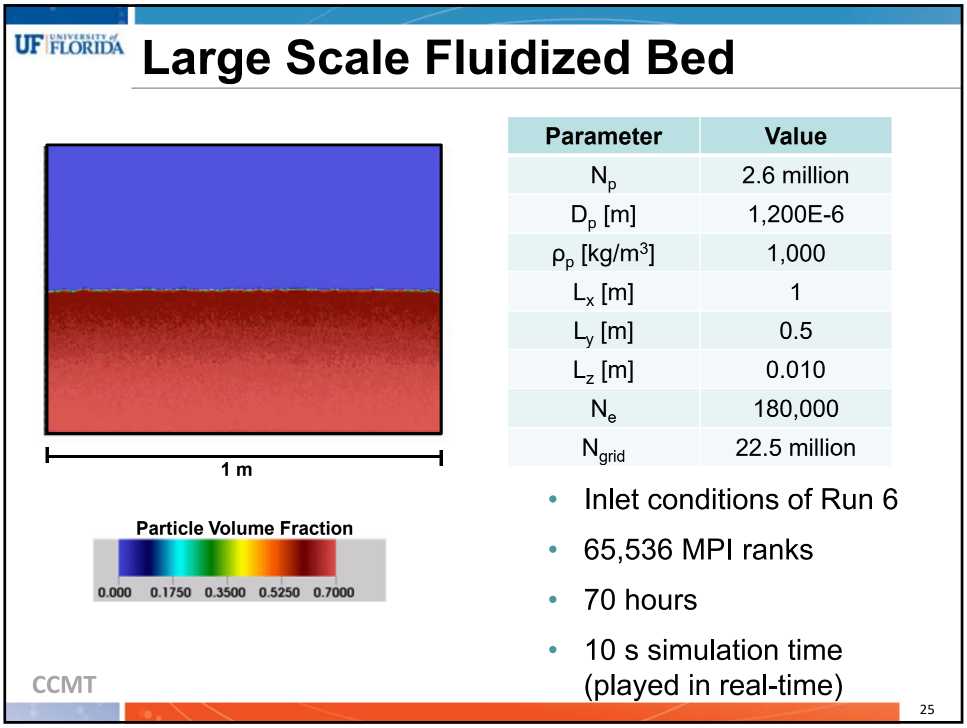

Large Scale Fluidized Bed

Parameter	Value
N_p	2.6 million
D_p [m]	1,200E-6
ρ_p [kg/m^3]	1,000
L_x [m]	1
L_y [m]	0.5
L_z [m]	0.010
N_e	180,000
N_{grid}	22.5 million

- Inlet conditions of Run 6
- 65,536 MPI ranks
- 70 hours
- 10 s simulation time (played in real-time)

CCMT

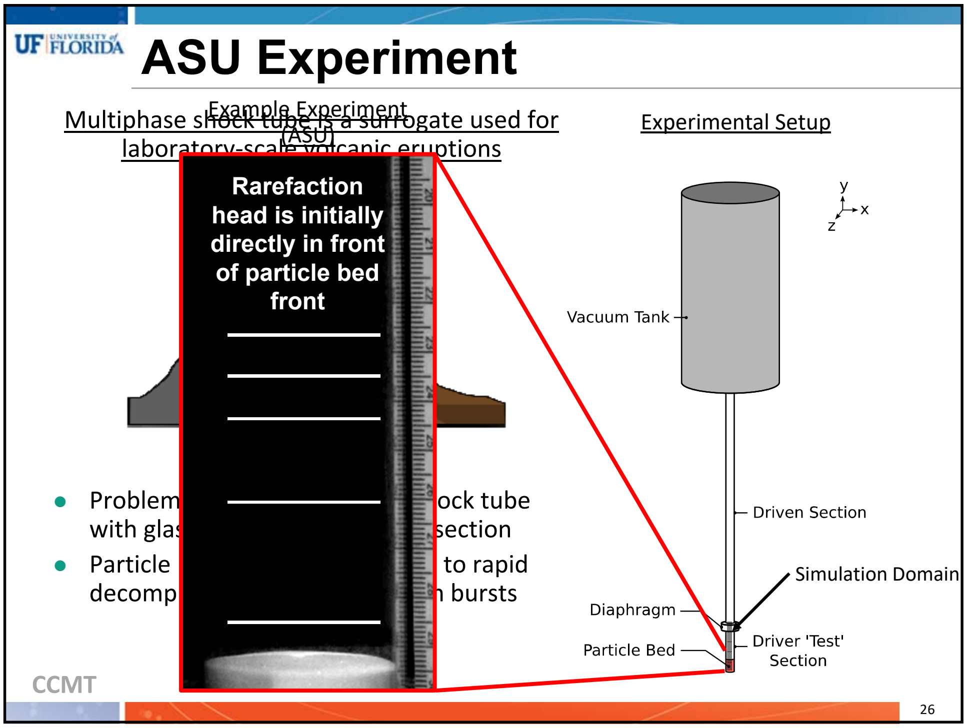

ASU Experiment

Multiphase shock tube is a surrogate used for laboratory-scale volcanic eruptions

Example Experiment (ASU)

Experimental Setup

- Problem … ock tube with glas … section
- Particle … to rapid decomp … n bursts

CCMT

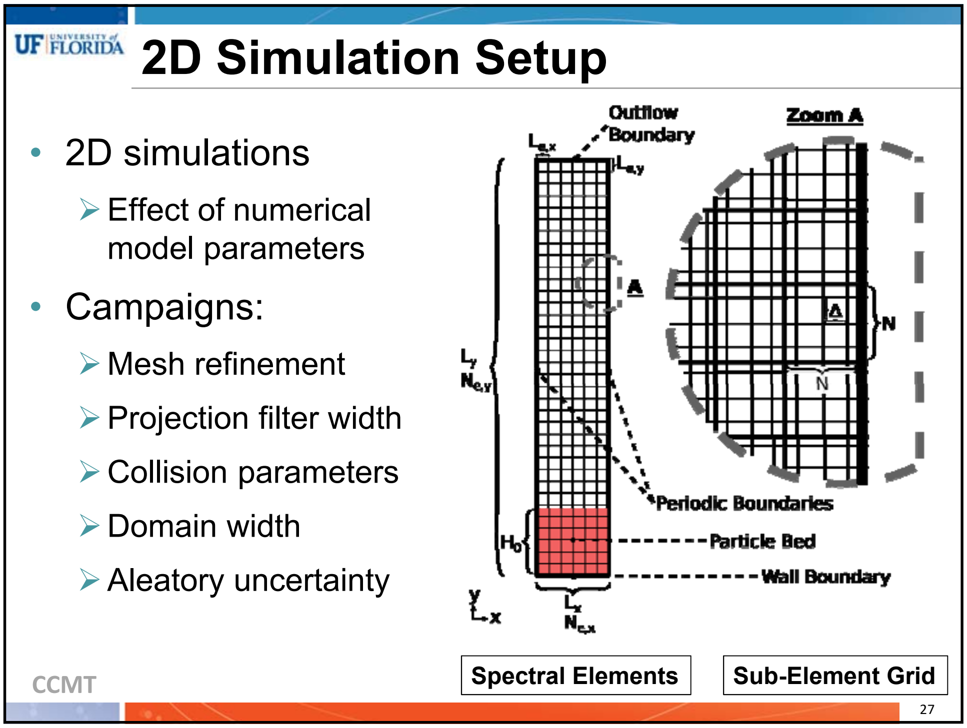

2D Simulation Setup

- 2D simulations
  - Effect of numerical model parameters
- Campaigns:
  - Mesh refinement
  - Projection filter width
  - Collision parameters
  - Domain width
  - Aleatory uncertainty

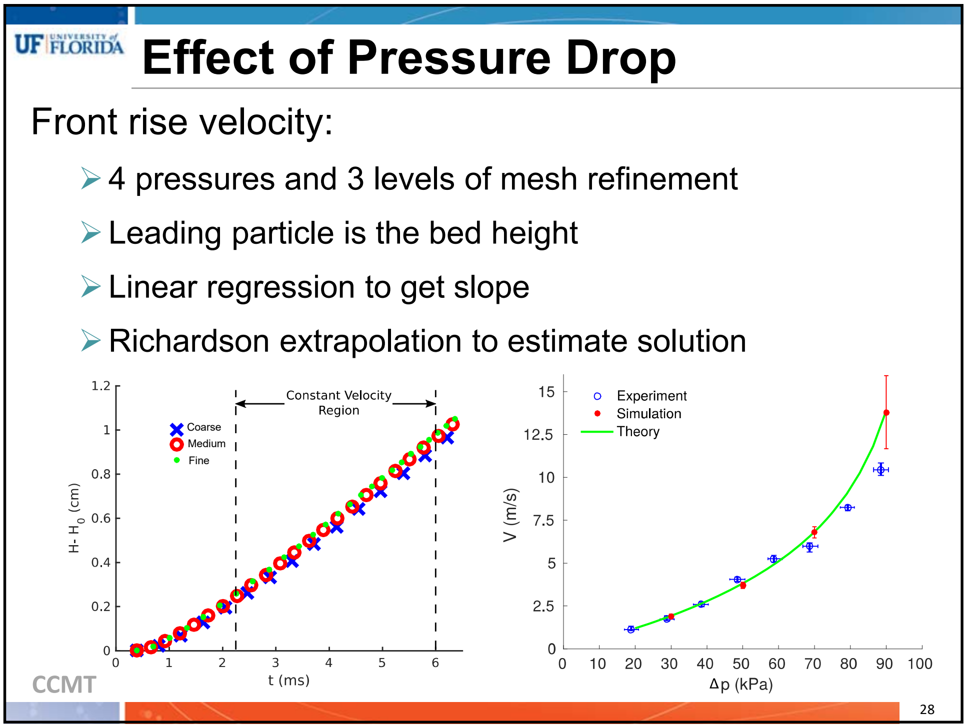

Effect of Pressure Drop

Front rise velocity:

- 4 pressures and 3 levels of mesh refinement
- Leading particle is the bed height
- Linear regression to get slope
- Richardson extrapolation to estimate solution

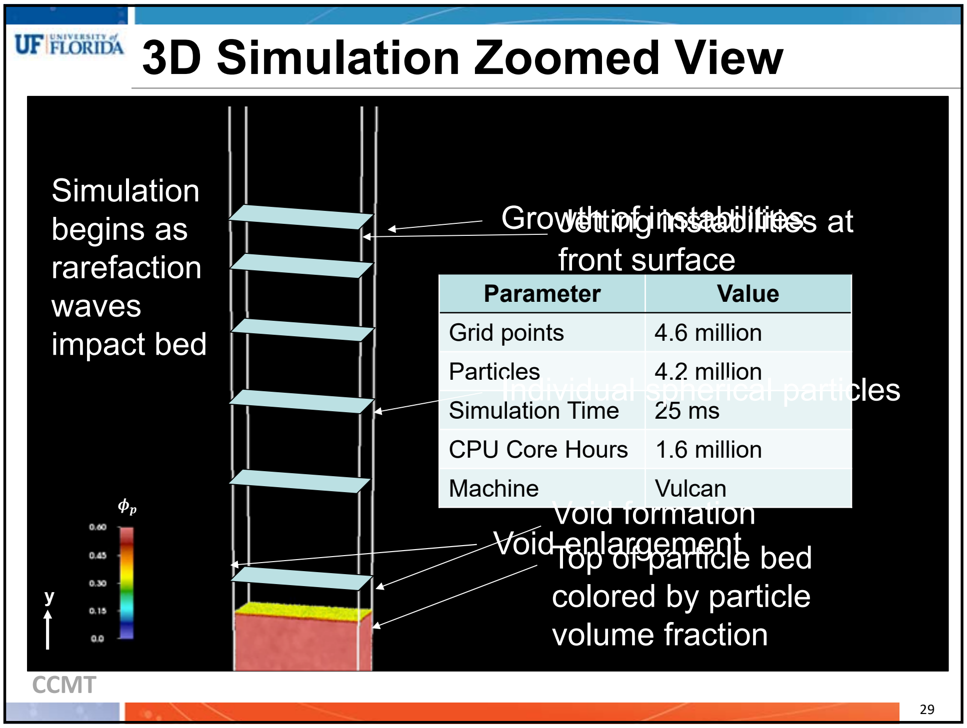

3D Simulation Zoomed View

Simulation begins as rarefaction waves impact bed

Growth of instabilities

Jetting instabilities at front surface

Parameter	Value
Grid points	4.6 million
Particles	4.2 million
Simulation Time	25 ms
CPU Core Hours	1.6 million
Machine	Vulcan

Individual spherical particles

Void formation

Void enlargement

Top of particle bed colored by particle volume fraction

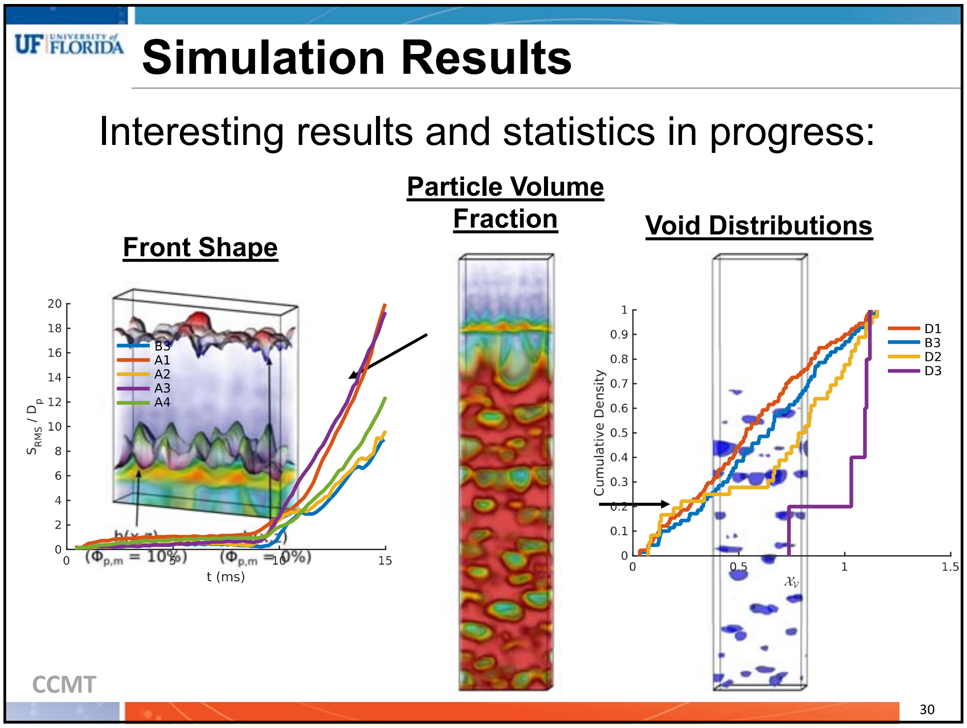

Simulation Results

Interesting results and statistics in progress:

Front Shape

Particle Volume Fraction

Void Distributions

Transition and code management

- CMT-nek workshop November 29, 2017
  - Training and introduction for CS, macroscale
- Host nek5000 User/Developer Meeting
  - Tampa, Florida April 17-18, 2018
- CMT-nek developer documentation
  - Doxygen driven
- Automated regression testing via travis.cl

Demonstration problem roadmap

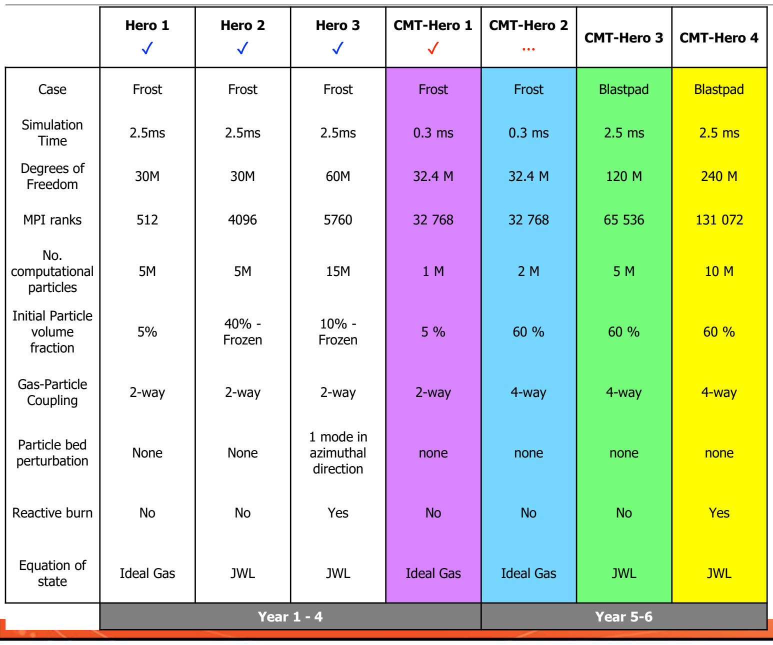

	Hero 1 ✓	Hero 2 ✓	Hero 3 ✓	CMT-Hero 1 ✓	CMT-Hero 2 …	CMT-Hero 3	CMT-Hero 4
Case	Frost	Frost	Frost	Frost	Frost	Blastpad	Blastpad
Simulation Time	2.5ms	2.5ms	2.5ms	0.3 ms	0.3 ms	2.5 ms	2.5 ms
Degrees of Freedom	30M	30M	60M	32.4 M	32.4 M	120 M	240 M
MPI ranks	512	4096	5760	32 768	32 768	65 536	131 072
No. computational particles	5M	5M	15M	1 M	2 M	5 M	10 M
Initial Particle volume fraction	5%	40% - Frozen	10% - Frozen	5 %	60 %	60 %	60 %
Gas-Particle Coupling	2-way	2-way	2-way	2-way	4-way	4-way	4-way
Particle bed perturbation	None	None	1 mode in azimuthal direction	none	none	none	none
Reactive burn	No	No	Yes	No	No	No	Yes
Equation of state	Ideal Gas	JWL	JWL	Ideal Gas	Ideal Gas	JWL	JWL
	Year 1 - 4				Year 5-6		

Transition plan: timeline

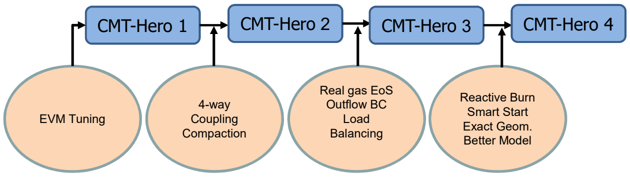

	Y6-Q1	Y6-Q2	Y6-Q3	Y6-Q4
CMT-Hero 3	🟧			
Outflow BC	🟦	🟦		
Load balancing	🟦			
Localized artificial viscosity		🟦	🟦	🟦
CMT-Hero 4	🟧	🟧	🟧	🟧
Reactive burn		🟦	🟦	
Smart start			🟦	🟦
Entropy-stable fluxes	🟦	🟦	🟦	

DGSEM is undergoing rapid change as its usage spreads!

- Better & cheaper ways to stabilize high-order methods

CCMT

Convective terms: past & future

Nonlinear flux functions **H**^{c}

$$U_j u_i + \phi_g p \delta_{ij}$$

$$\underset{\mathbb{P}^{N-1}}{\mathbf{v}^\intercal}\mathbf{D}_k^\intercal\mathbf{B}\left[\mathbf{r}_{k,\mathbf{x}_i}\right]\underset{\mathbb{P}^{N-1}}{[U_j]}\underset{?}{[u_i]}$$

- Integrand has products of polynomial & rational functions (and thermodynamic properties)
- Most stability proofs for DG depend on exact integration, but quadrature on rational functions is not exact.

But it turns out that $\mathcal{D}$ has the summation-by-parts (SBP) property

$$(\mathbf{B}\mathcal{D})^\intercal = \mathrm{diag}(-1,0,\ldots,0,1) - \mathbf{B}\mathcal{D}$$

Which means it discretely guarantees the weighted-residual DG statement for a wider range of functions than the quadrature weights in **B** can integrate exactly

$$\int_{\Omega_e} v\nabla\cdot\mathbf{H}\mathrm{d}V = \int_{\partial\Omega_e} v\mathbf{H}^*\cdot\mathbf{n}\mathrm{d}A - \int_{\Omega_e}(\nabla v)\cdot\mathbf{H}\mathrm{d}V$$

Provided **H**^{c} is written in a very special way…

Gassner, Winters & Kopriva (2016) "Split form nodal DG schemes with SBP for Euler equations" *JCP* **327**:39-66

Carpenter *et al* (2014) *SIAM J. Sci. Comput.* **36**(5):B835-B867

CCMT

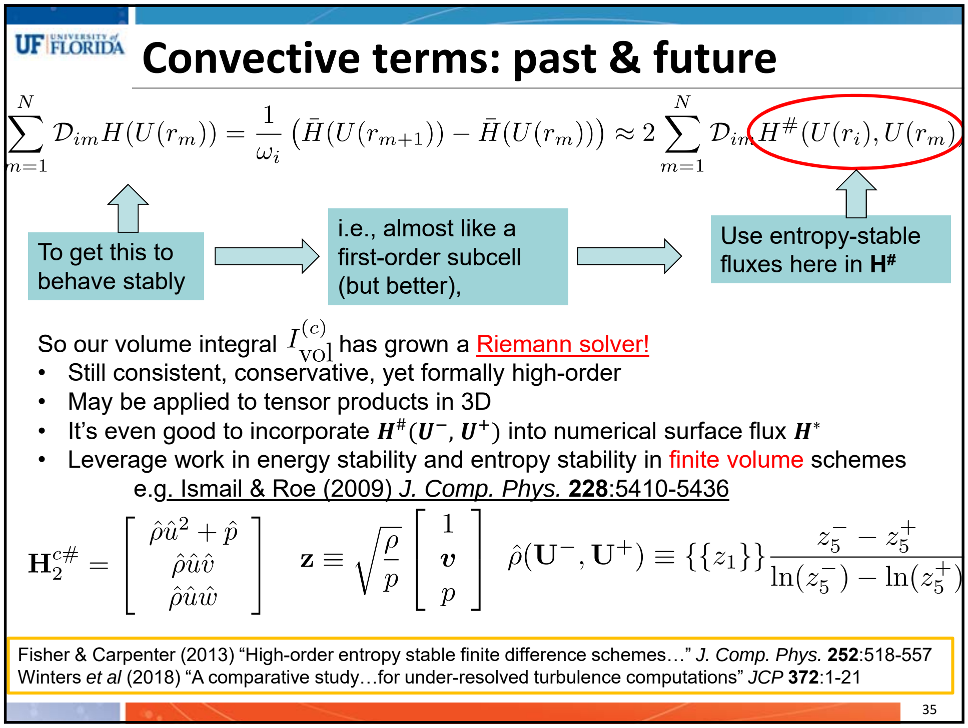

Convective terms: past & future

$$\sum_{m=1}^{N} \mathcal{D}_{im} H(U(r_m)) = \frac{1}{\omega_i}\left(\bar{H}(U(r_{m+1})) - \bar{H}(U(r_m))\right) \approx 2\sum_{m=1}^{N} \mathcal{D}_{im} H^{\#}(U(r_i), U(r_m))$$

To get this to behave stably

i.e., almost like a first-order subcell (but better),

Use entropy-stable fluxes here in **H#**

So our volume integral $I_{\text{vol}}^{(c)}$ has grown a Riemann solver!

- Still consistent, conservative, yet formally high-order
- May be applied to tensor products in 3D
- It's even good to incorporate $\boldsymbol{H}^{\#}(\boldsymbol{U}^{-}, \boldsymbol{U}^{+})$ into numerical surface flux $\boldsymbol{H}^{*}$
- Leverage work in energy stability and entropy stability in finite volume schemes

e.g. Ismail & Roe (2009) *J. Comp. Phys.* **228**:5410-5436

$$\mathbf{H}_2^{c\#} = \begin{bmatrix} \hat{\rho}\hat{u}^2 + \hat{p} \\ \hat{\rho}\hat{u}\hat{v} \\ \hat{\rho}\hat{u}\hat{w} \end{bmatrix} \qquad \mathbf{z} \equiv \sqrt{\frac{\rho}{p}} \begin{bmatrix} 1 \\ \boldsymbol{v} \\ p \end{bmatrix} \qquad \hat{\rho}(\mathbf{U}^{-}, \mathbf{U}^{+}) \equiv \{\{z_1\}\} \frac{z_5^{-} - z_5^{+}}{\ln(z_5^{-}) - \ln(z_5^{+})}$$

Fisher & Carpenter (2013) "High-order entropy stable finite difference schemes…" *J. Comp. Phys.* **252**:518-557
Winters *et al* (2018) "A comparative study…for under-resolved turbulence computations" *JCP* **372**:1-21

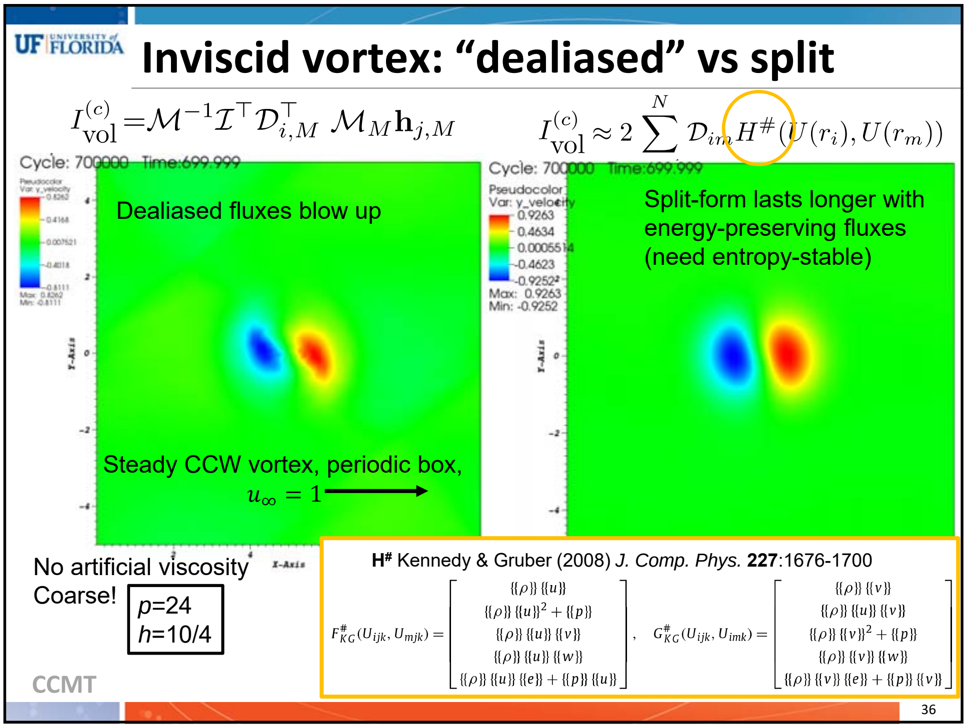

Inviscid vortex: "dealiased" vs split

$$I_{\text{vol}}^{(c)} = \mathcal{M}^{-1}\mathcal{I}^{\top}\mathcal{D}_{i,M}^{\top}\ \mathcal{M}_M \mathbf{h}_{j,M} \qquad I_{\text{vol}}^{(c)} \approx 2\sum^{N} \mathcal{D}_{im} H^{\#}(U(r_i), U(r_m))$$

Dealiased fluxes blow up

Split-form lasts longer with energy-preserving fluxes (need entropy-stable)

Steady CCW vortex, periodic box, $u_\infty = 1 \longrightarrow$

No artificial viscosity

Coarse! p=24 h=10/4

H# Kennedy & Gruber (2008) *J. Comp. Phys.* **227**:1676-1700

$$F_{KG}^{\#}(U_{ijk}, U_{mjk}) = \begin{bmatrix} \{\{\rho\}\}\,\{\{u\}\} \\ \{\{\rho\}\}\,\{\{u\}\}^2 + \{\{p\}\} \\ \{\{\rho\}\}\,\{\{u\}\}\,\{\{v\}\} \\ \{\{\rho\}\}\,\{\{u\}\}\,\{\{w\}\} \\ \{\{\rho\}\}\,\{\{u\}\}\,\{\{e\}\} + \{\{p\}\}\,\{\{u\}\} \end{bmatrix}, \quad G_{KG}^{\#}(U_{ijk}, U_{imk}) = \begin{bmatrix} \{\{\rho\}\}\,\{\{v\}\} \\ \{\{\rho\}\}\,\{\{u\}\}\,\{\{v\}\} \\ \{\{\rho\}\}\,\{\{v\}\}^2 + \{\{p\}\} \\ \{\{\rho\}\}\,\{\{v\}\}\,\{\{w\}\} \\ \{\{\rho\}\}\,\{\{v\}\}\,\{\{e\}\} + \{\{p\}\}\,\{\{v\}\} \end{bmatrix}$$

CCMT

UF UNIVERSITY of FLORIDA

Summary and Future Work

- CMT-nek is filling RocFlu's shoes and beyond
  - ➢ Compressible multiphase flow, scalable
  - ➢ Production simulation of microscale, mesoscale and macroscale flows
  - ➢ Behavioral emulation companion suite
- Future work (new capability)
  - ➢ Nonreflecting boundary conditions
  - ➢ Reactive burn initial conditions
  - ➢ Localizing artificial viscosity
    - ➢ Shock detectors
    - ➢ Improved numerical fluxes for discontinuous AV (Weighted interior penalty, Zunino (2009) *SIAM J. Sci. Comput.* **38**:99-126)
  - ➢ Entropy stability (Carpenter *et al* (2014) *SIAM J. Sci. Comput.* **36**)
    - ➢ Entropy-stable fluxes for multiphase systems, non-ideal EOS

CCMT

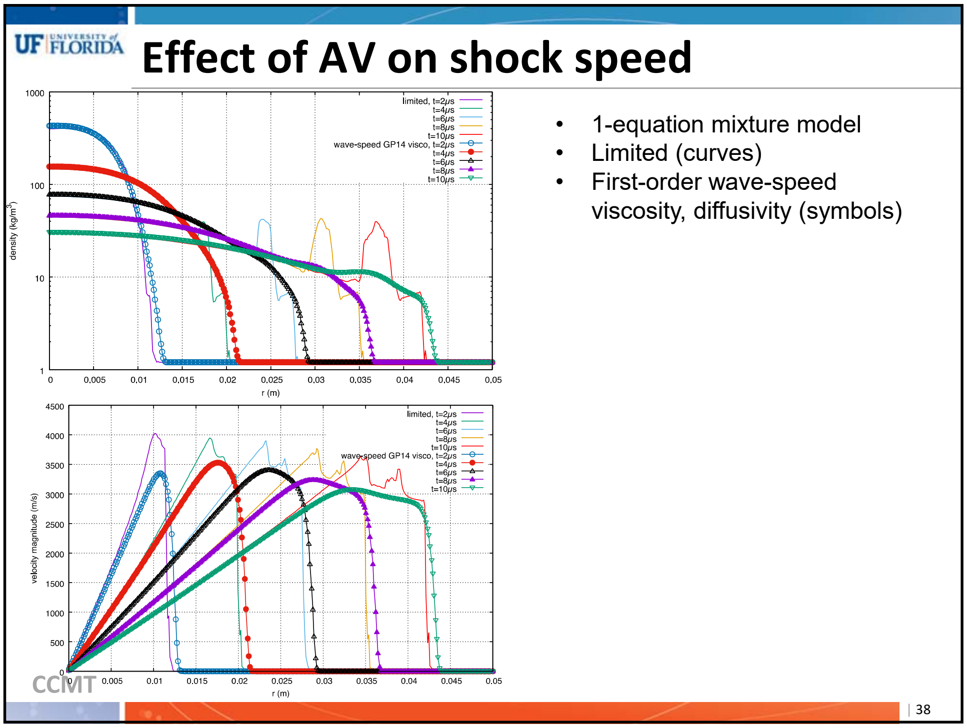

- 1-equation mixture model
- Limited (curves)
- First-order wave-speed viscosity, diffusivity (symbols)

CCMT

Do you have any questions?

UF UNIVERSITY of FLORIDA

NNSA

CCMT

Uncertainty Budget

Validation and Uncertainty Reduction

Chanyoung Park, Raphael (Rafi) T. Haftka and Nam-Ho Kim

Department of Mechanical & Aerospace Engineering

Uncertainty Budget Team


Raphael Haftka, PhD


Nam-Ho Kim, PhD


Sangjune Bae


Samaun Nili


Chanyoung Park, PhD


Kyle Hughes, PhD

- PIs: Raphael Haftka and Nam-Ho Kim
- Research Scientist: Chanyoung Park (Joining **Caterpillar** at Champaign)
- Students
  - Samaun Nili (Mesoscale force model potential errors)
  - Sangjune Bae (Joining **CCMT as postdoc**)
  - Kyle Hughes (Graduated; **Postdoc** at **LANL**)
  - Giselle Fernandez (Graduated; **Postdoc at LANL** or **SNL**)


Giselle Fernandez

CCMT

UQ and Validation of Demonstration Problem

Forensic UQ for designing the experiment

UQ of experiments to obtain simulation inputs and quantities of interest (QoI)

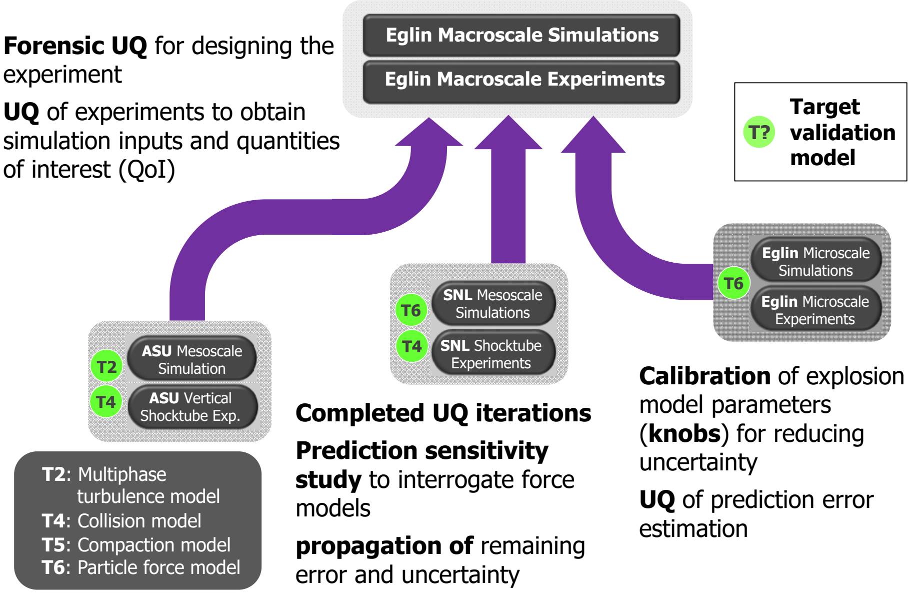

Eglin Macroscale Simulations

Eglin Macroscale Experiments

T? **Target validation model**

T2 **ASU** Mesoscale Simulation

T4 **ASU** Vertical Shocktube Exp.

T6 **SNL** Mesoscale Simulations

T4 **SNL** Shocktube Experiments

T6 **Eglin** Microscale Simulations

Eglin Microscale Experiments

T2: Multiphase turbulence model
T4: Collision model
T5: Compaction model
T6: Particle force model

Completed UQ iterations

Prediction sensitivity study to interrogate force models

propagation of remaining error and uncertainty

Calibration of explosion model parameters (**knobs**) for reducing uncertainty

UQ of prediction error estimation

Outline

- Eglin microscale
  - **Calibration** and **UQ** of reactive burn model parameters (knobs)
  - **UQ** of prediction error estimation of particle force model
- UQ of Eglin macroscale
  - **Propagation** of error and uncertainty from Sandia meso- to Eglin macroscale based on sensitivity study
  - **Forensic UQ** based on past experiments

Microscale UQ

UQ of simulation parameters and UR through calibration

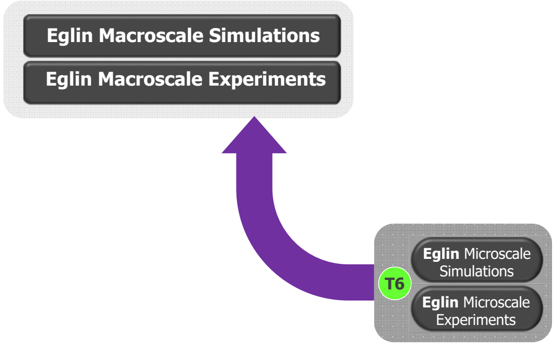

Calibration of explosion model parameters (**knobs**) for reducing uncertainty

UQ of prediction error estimation

Eglin Microscale Experiment

- Calibration of reactive burn model parameters
- UQ of JWL parameters
- Assessment of particle force model
- Particle position (PP) is quantity of interest (QoI)

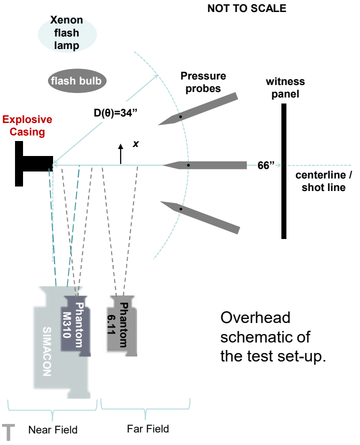

Overhead schematic of the test set-up.

Explosive Casing

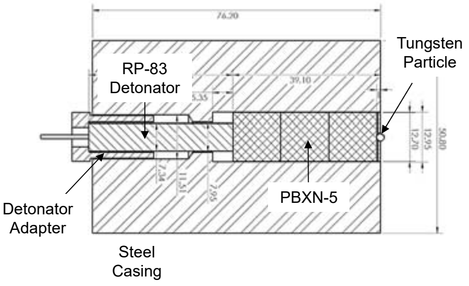

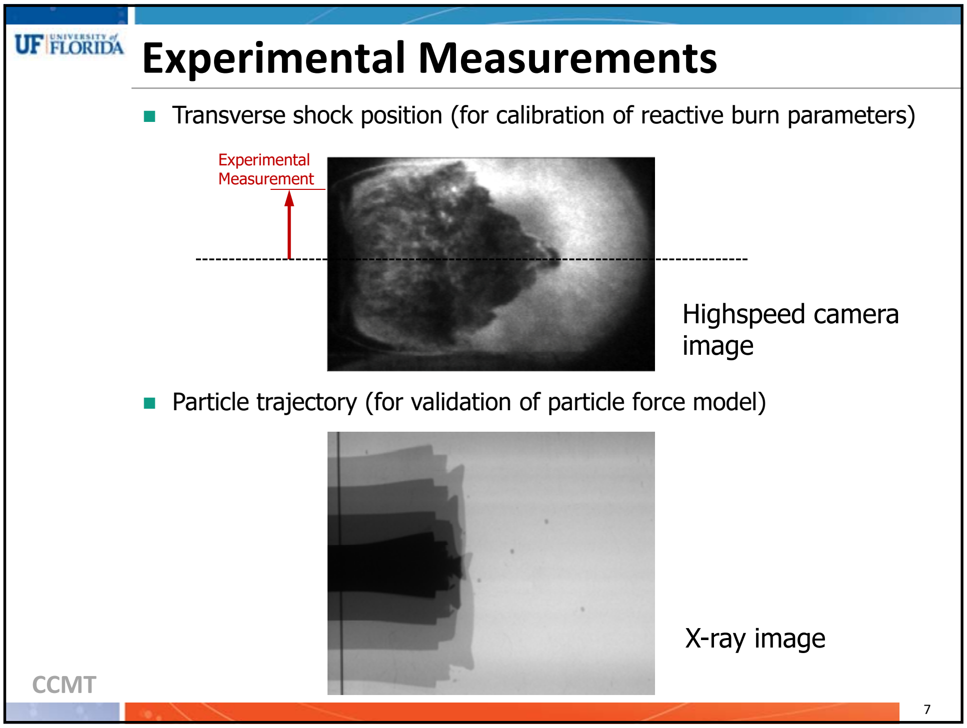

Experimental Measurements

- Transverse shock position (for calibration of reactive burn parameters)

Highspeed camera image

- Particle trajectory (for validation of particle force model)

X-ray image

Uncertain Parameters

Parameter	Quantity	Method
U1: Explosive density	mean = 1795 kg/m^3, std = 2.9 kg/m^3, Bi-modal	Derived
U2: RZ thickness	[0.365, 1.4] mm	Expert opinion
U3: Heat release Q	[11.6, 14.2] MJ/kg	Literature and opinion
U4: Particle diameter	mean = 2.0156 mm, std = 0.0073 mm, Weibull	Direct measurement Micrometer, 52 samples
U5: Particle density	mean = 15540 kg/m^3, std = 250 kg/m^3, Normal	Pycnometer Gas Pycnometer, 12 samples
U6: Initial radial position	[0, 0.254] mm	
U7-11: JWL	A, B, w, R_1, R_2	Literature and opinion
…	…	…

- Listed uncertainties including uncertainties in simulation parameters (knob)
- **Kyle** actively involved in measurements and uncertainty estimation (forensic UQ)

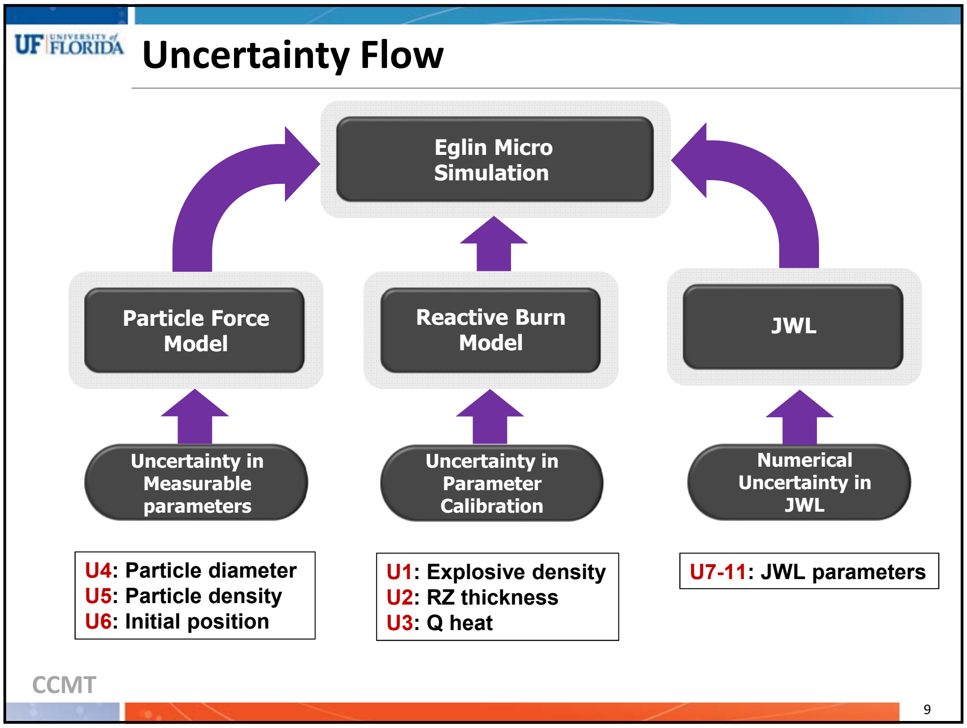

Uncertainty Flow

Eglin Micro Simulation

Particle Force Model

Reactive Burn Model

JWL

Uncertainty in Measurable parameters

Uncertainty in Parameter Calibration

Numerical Uncertainty in JWL

U4: Particle diameter
U5: Particle density
U6: Initial position

U1: Explosive density
U2: RZ thickness
U3: Q heat

U7-11: JWL parameters

CCMT

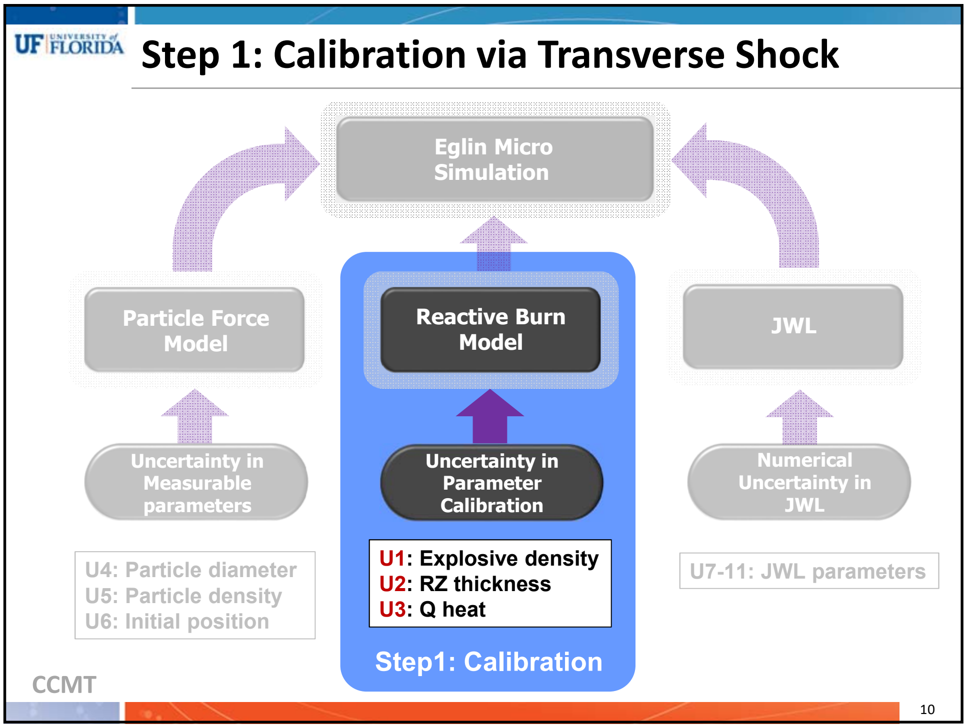

Step 1: Calibration via Transverse Shock

Eglin Micro Simulation

Particle Force Model

Reactive Burn Model

JWL

Uncertainty in Measurable parameters

Uncertainty in Parameter Calibration

Numerical Uncertainty in JWL

U4: Particle diameter
U5: Particle density
U6: Initial position

U1: Explosive density
U2: RZ thickness
U3: Q heat

U7-11: JWL parameters

Step1: Calibration

CCMT

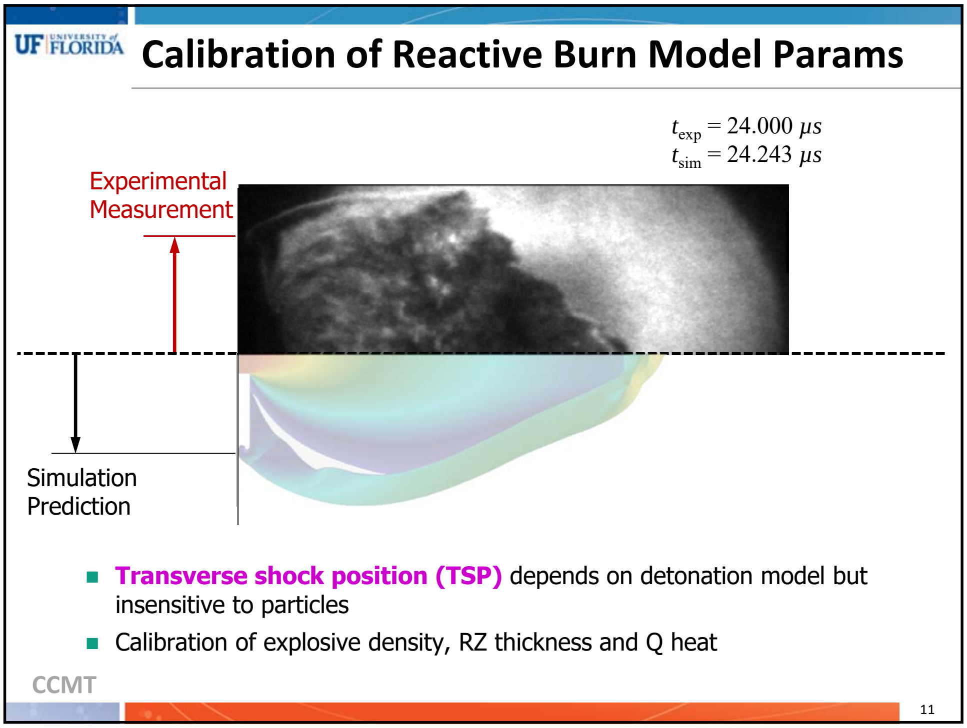

Calibration of Reactive Burn Model Params

$t_{\text{exp}} = 24.000\ \mu s$
$t_{\text{sim}} = 24.243\ \mu s$

- **Transverse shock position (TSP)** depends on detonation model but insensitive to particles
- Calibration of explosive density, RZ thickness and Q heat

CCMT

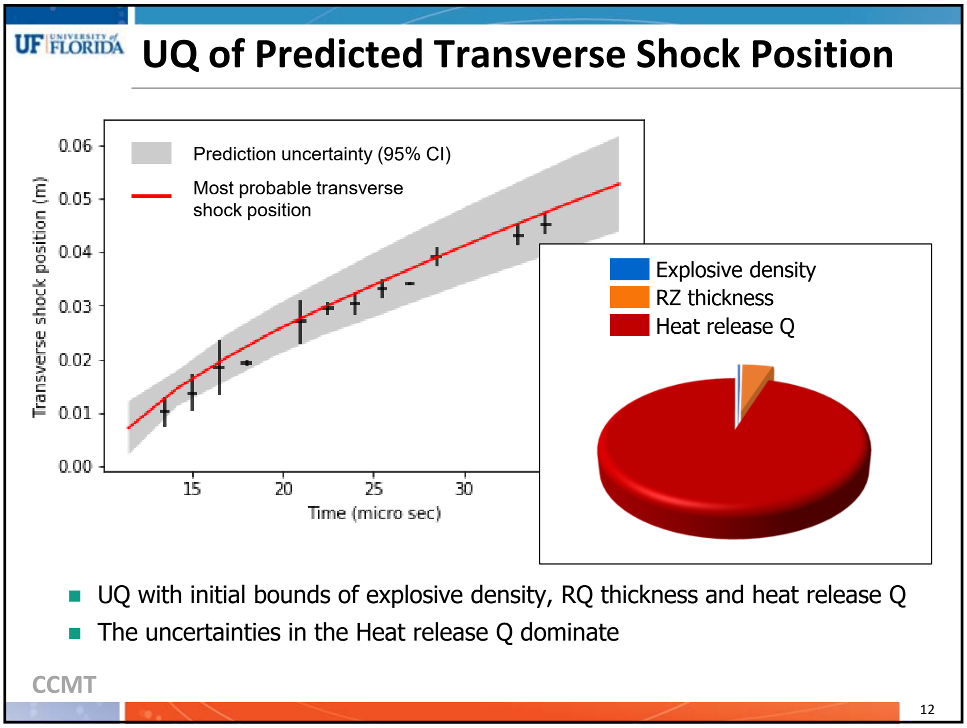

UQ of Predicted Transverse Shock Position

- UQ with initial bounds of explosive density, RQ thickness and heat release Q
- The uncertainties in the Heat release Q dominate

CCMT

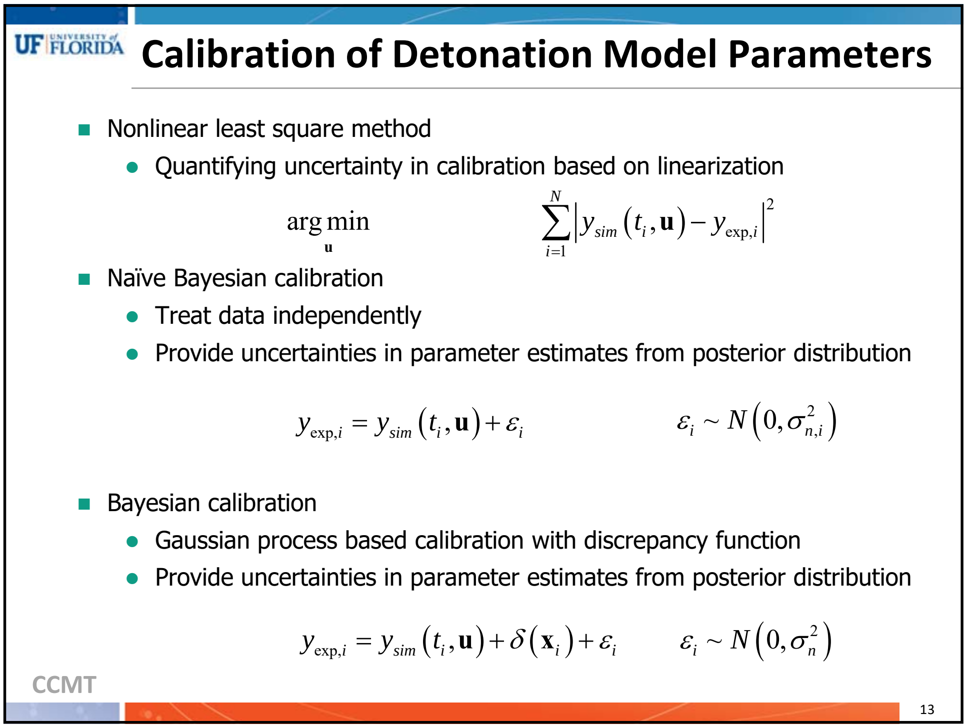

Calibration of Detonation Model Parameters

- Nonlinear least square method
  - Quantifying uncertainty in calibration based on linearization

$$\underset{\mathbf{u}}{\arg\min} \qquad \sum_{i=1}^{N} \left| y_{sim}\left(t_i,\mathbf{u}\right) - y_{\exp,i} \right|^2$$

- Naïve Bayesian calibration
  - Treat data independently
  - Provide uncertainties in parameter estimates from posterior distribution

$$y_{\exp,i} = y_{sim}\left(t_i,\mathbf{u}\right) + \varepsilon_i \qquad \varepsilon_i \sim N\left(0,\sigma_{n,i}^2\right)$$

- Bayesian calibration
  - Gaussian process based calibration with discrepancy function
  - Provide uncertainties in parameter estimates from posterior distribution

$$y_{\exp,i} = y_{sim}\left(t_i,\mathbf{u}\right) + \delta\left(\mathbf{x}_i\right) + \varepsilon_i \qquad \varepsilon_i \sim N\left(0,\sigma_n^2\right)$$

CCMT

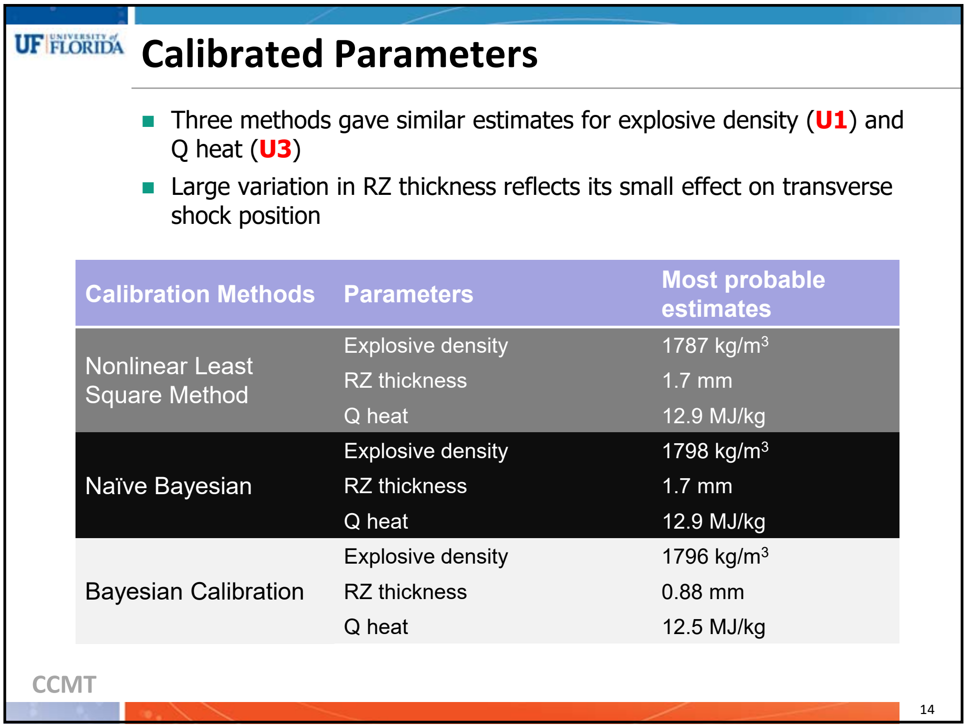

Calibrated Parameters

- Three methods gave similar estimates for explosive density (**U1**) and Q heat (**U3**)
- Large variation in RZ thickness reflects its small effect on transverse shock position

Calibration Methods	Parameters	Most probable estimates
Nonlinear Least Square Method	Explosive density	1787 kg/m^3
	RZ thickness	1.7 mm
	Q heat	12.9 MJ/kg
Naïve Bayesian	Explosive density	1798 kg/m^3
	RZ thickness	1.7 mm
	Q heat	12.9 MJ/kg
Bayesian Calibration	Explosive density	1796 kg/m^3
	RZ thickness	0.88 mm
	Q heat	12.5 MJ/kg

CCMT

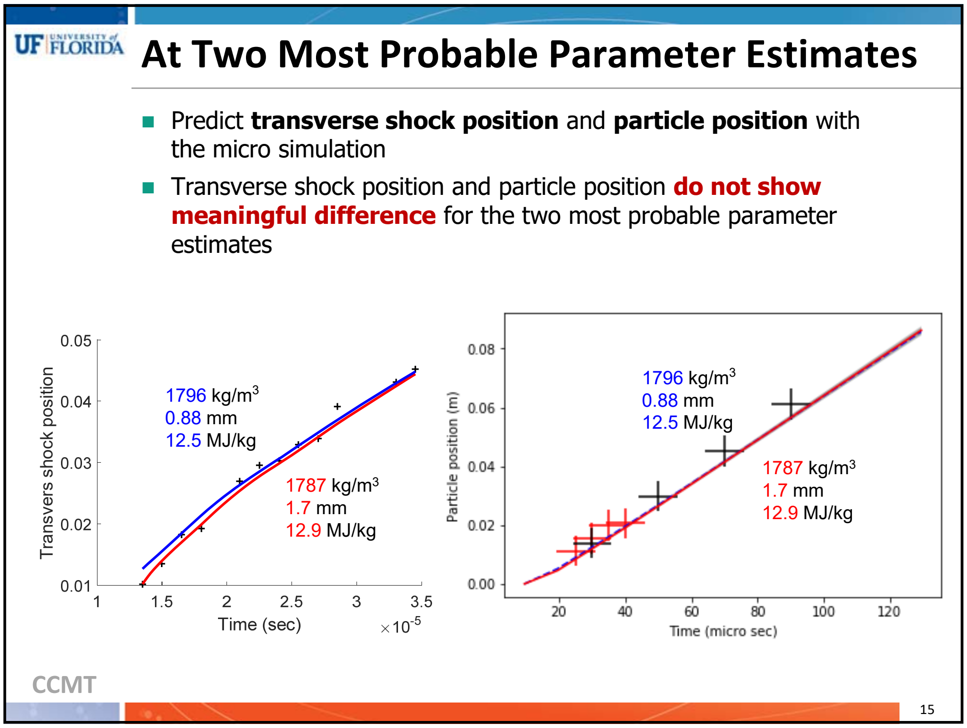

At Two Most Probable Parameter Estimates

- Predict **transverse shock position** and **particle position** with the micro simulation
- Transverse shock position and particle position **do not show meaningful difference** for the two most probable parameter estimates

CCMT

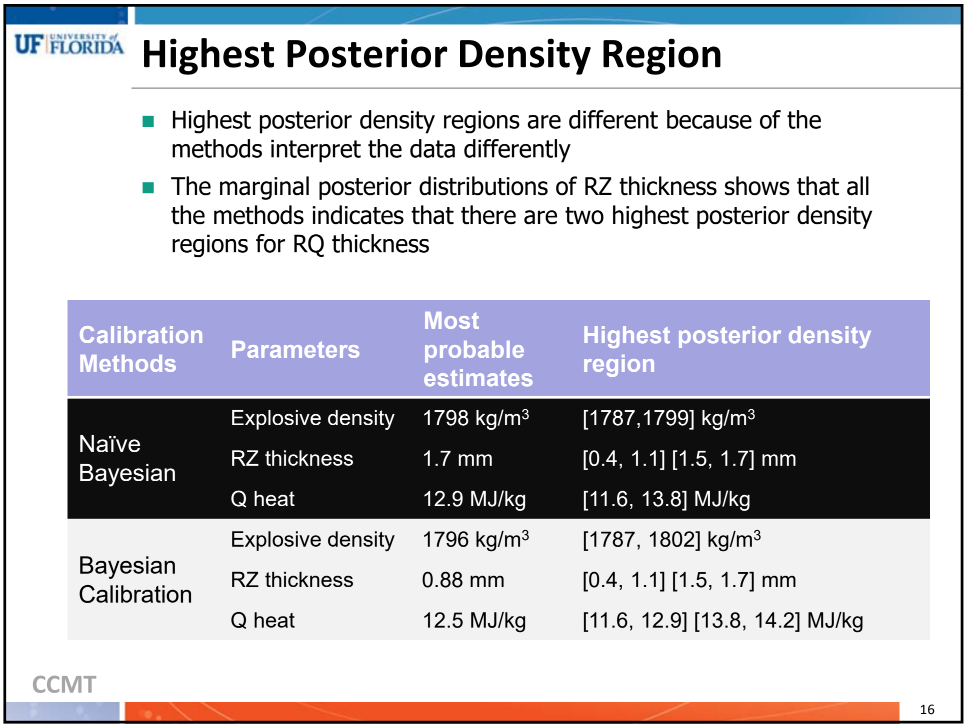

Highest Posterior Density Region

- Highest posterior density regions are different because of the methods interpret the data differently
- The marginal posterior distributions of RZ thickness shows that all the methods indicates that there are two highest posterior density regions for RQ thickness

Calibration Methods	Parameters	Most probable estimates	Highest posterior density region
Naïve Bayesian	Explosive density	1798 kg/m^3	[1787,1799] kg/m^3
	RZ thickness	1.7 mm	[0.4, 1.1] [1.5, 1.7] mm
	Q heat	12.9 MJ/kg	[11.6, 13.8] MJ/kg
Bayesian Calibration	Explosive density	1796 kg/m^3	[1787, 1802] kg/m^3
	RZ thickness	0.88 mm	[0.4, 1.1] [1.5, 1.7] mm
	Q heat	12.5 MJ/kg	[11.6, 12.9] [13.8, 14.2] MJ/kg

CCMT

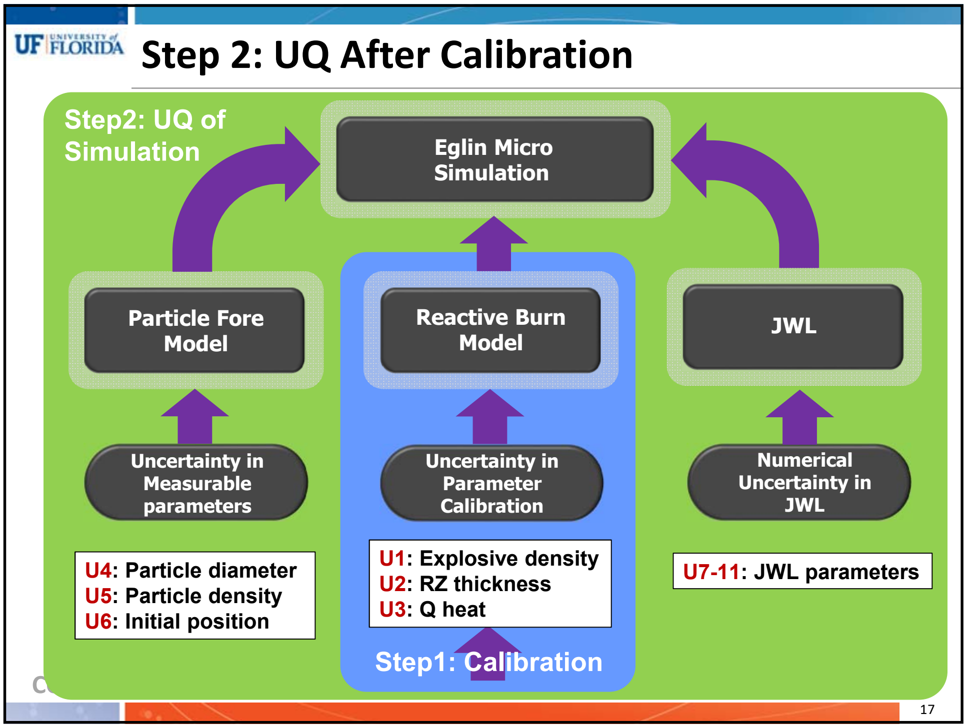

Step 2: UQ After Calibration

Step2: UQ of Simulation

Eglin Micro Simulation

Particle Fore Model

Reactive Burn Model

JWL

Uncertainty in Measurable parameters

Uncertainty in Parameter Calibration

Numerical Uncertainty in JWL

U4: Particle diameter
U5: Particle density
U6: Initial position

U1: Explosive density
U2: RZ thickness
U3: Q heat

U7-11: JWL parameters

Step1: Calibration

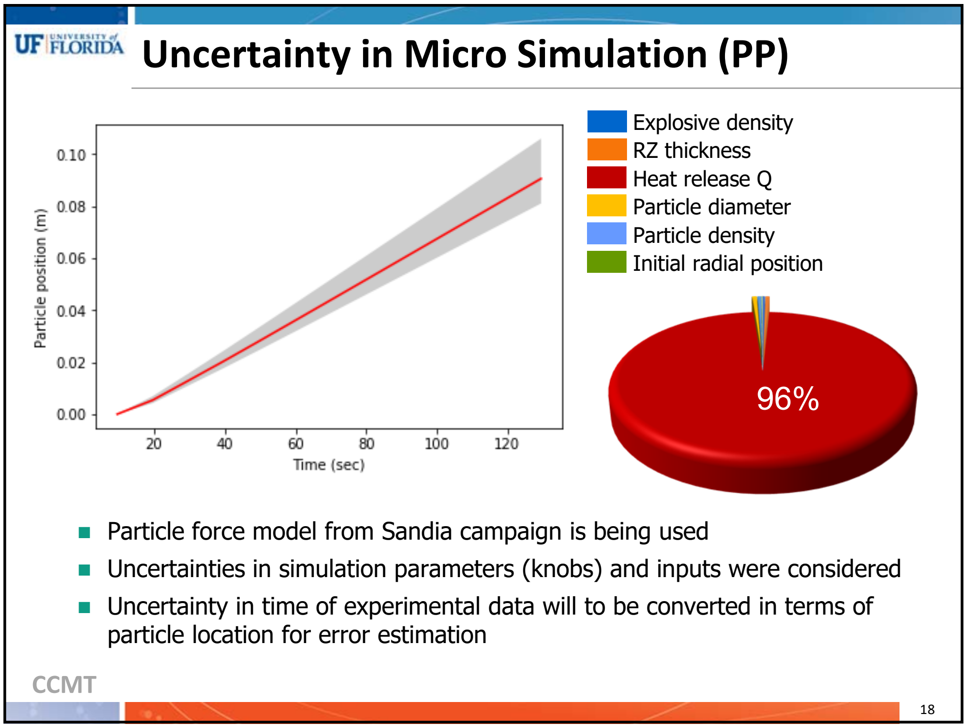

Uncertainty in Micro Simulation (PP)

- Particle force model from Sandia campaign is being used
- Uncertainties in simulation parameters (knobs) and inputs were considered
- Uncertainty in time of experimental data will to be converted in terms of particle location for error estimation

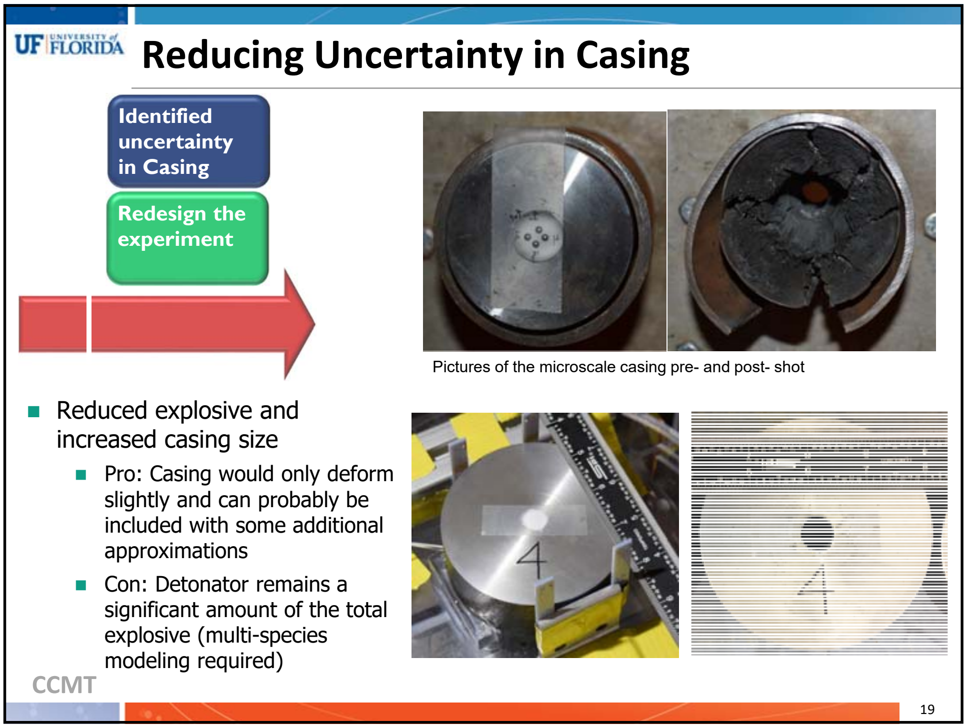

Reducing Uncertainty in Casing

Identified uncertainty in Casing

Redesign the experiment

Pictures of the microscale casing pre- and post- shot

- Reduced explosive and increased casing size
  - Pro: Casing would only deform slightly and can probably be included with some additional approximations
  - Con: Detonator remains a significant amount of the total explosive (multi-species modeling required)

CCMT

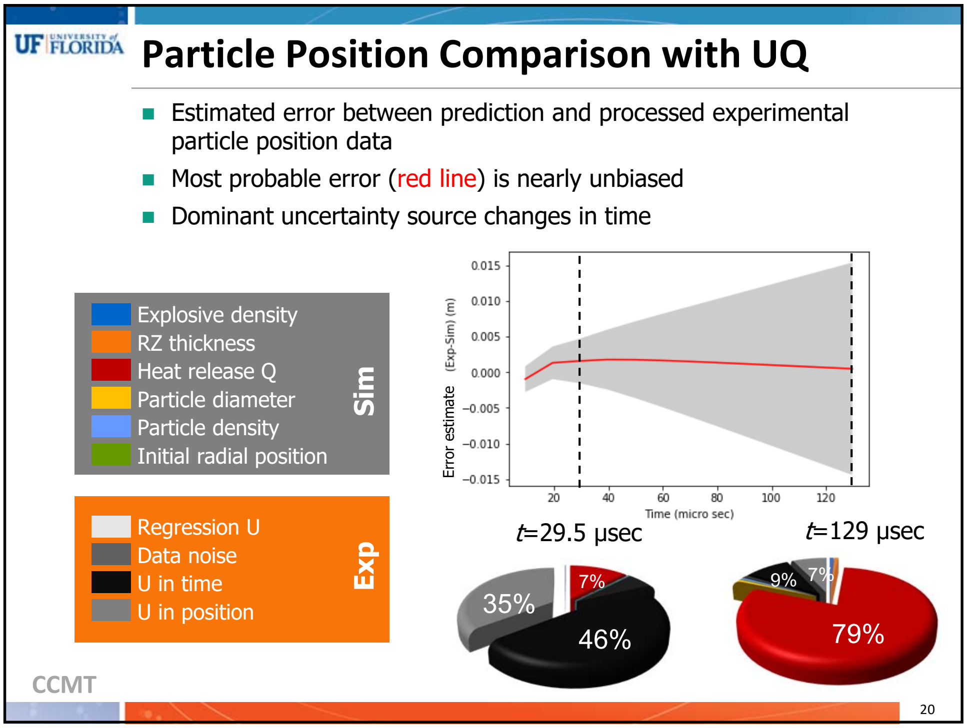

Particle Position Comparison with UQ

- Estimated error between prediction and processed experimental particle position data
- Most probable error (red line) is nearly unbiased
- Dominant uncertainty source changes in time

Sim

🟦 Explosive density
🟧 RZ thickness
🟥 Heat release Q
🟨 Particle diameter
🟦 Particle density
🟩 Initial radial position

Exp

⬜ Regression U
⬛ Data noise
⬛ U in time
⬛ U in position

CCMT

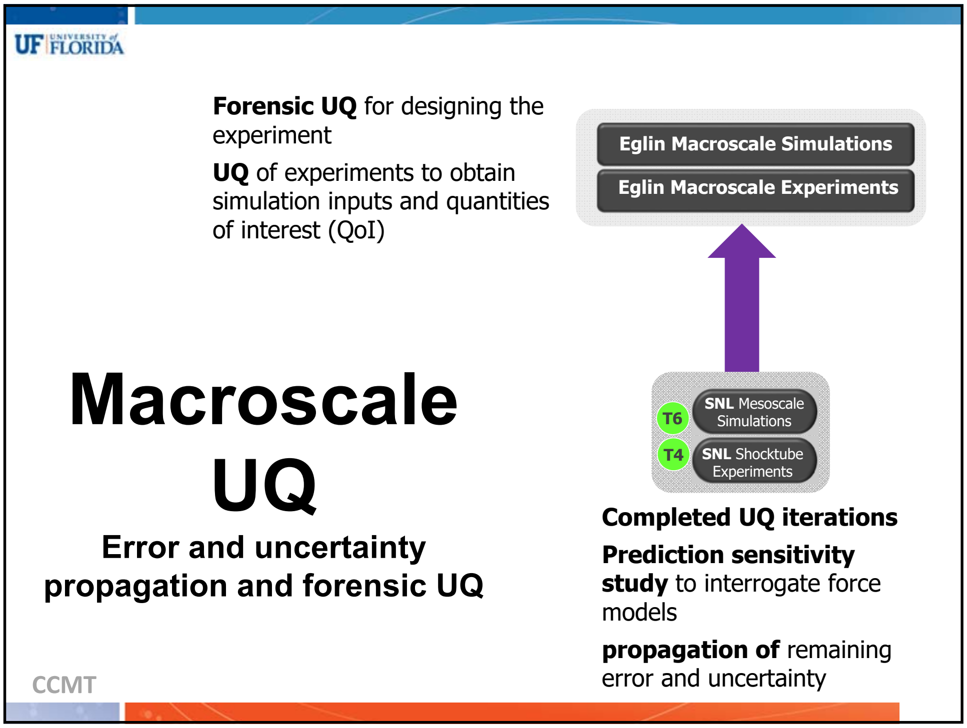

Forensic UQ for designing the experiment

UQ of experiments to obtain simulation inputs and quantities of interest (QoI)

Macroscale UQ

Error and uncertainty propagation and forensic UQ

Eglin Macroscale Simulations

Eglin Macroscale Experiments

T6 **SNL** Mesoscale Simulations

T4 **SNL** Shocktube Experiments

Completed UQ iterations

Prediction sensitivity study to interrogate force models

propagation of remaining error and uncertainty

CCMT

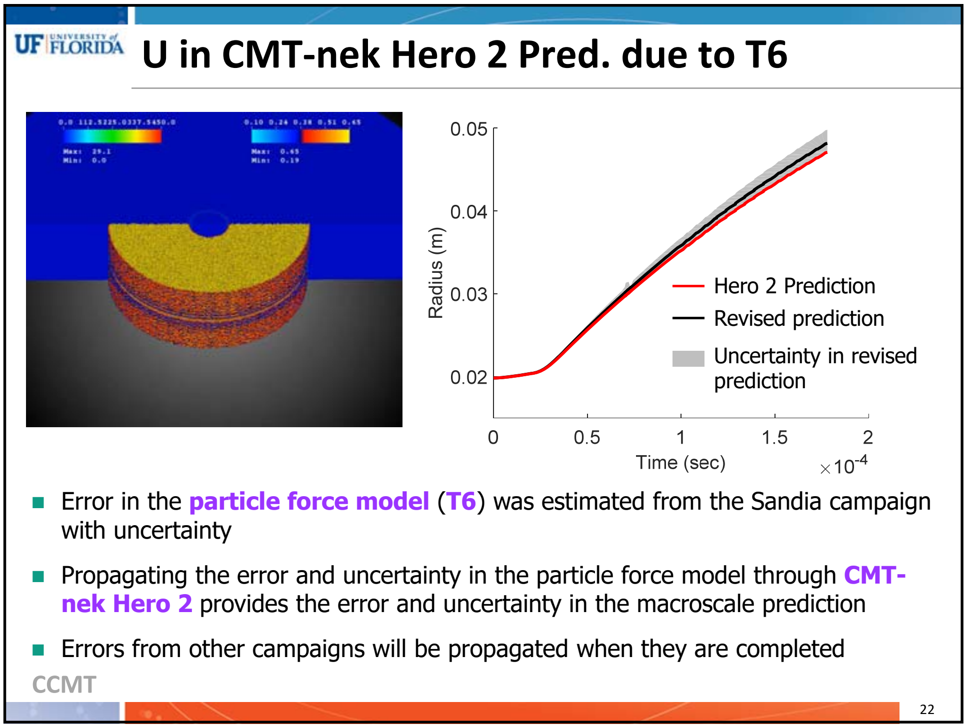

U in CMT-nek Hero 2 Pred. due to T6

- Error in the **particle force model** (**T6**) was estimated from the Sandia campaign with uncertainty
- Propagating the error and uncertainty in the particle force model through **CMT-nek Hero 2** provides the error and uncertainty in the macroscale prediction
- Errors from other campaigns will be propagated when they are completed

CCMT

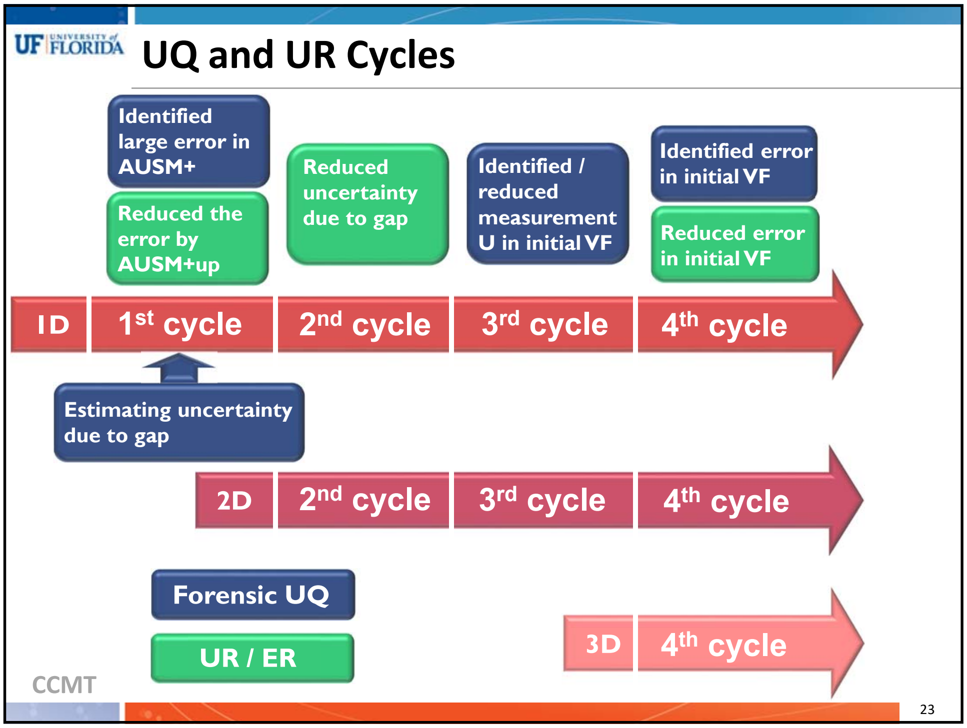

UQ and UR Cycles

Identified large error in AUSM+

Reduced the error by AUSM+up

Reduced uncertainty due to gap

Identified / reduced measurement U in initial VF

Identified error in initial VF

Reduced error in initial VF

1D | **1st cycle** | **2nd cycle** | **3rd cycle** | **4th cycle**

Estimating uncertainty due to gap

2D | **2nd cycle** | **3rd cycle** | **4th cycle**

Forensic UQ

UR / ER

3D | **4th cycle**

CCMT

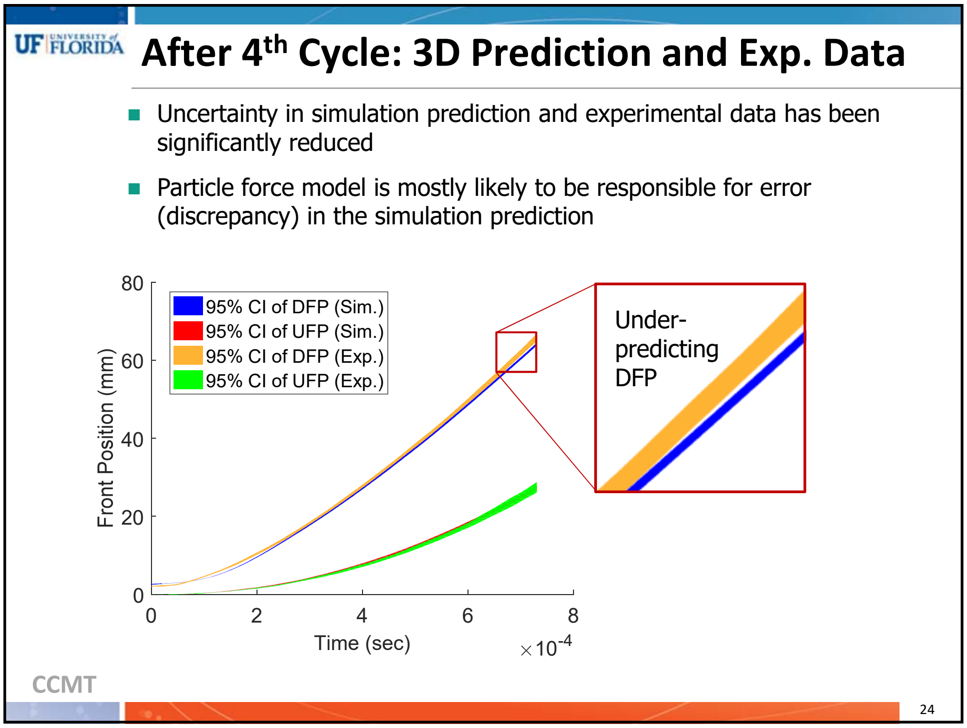

After 4th Cycle: 3D Prediction and Exp. Data

- Uncertainty in simulation prediction and experimental data has been significantly reduced
- Particle force model is mostly likely to be responsible for error (discrepancy) in the simulation prediction

CCMT

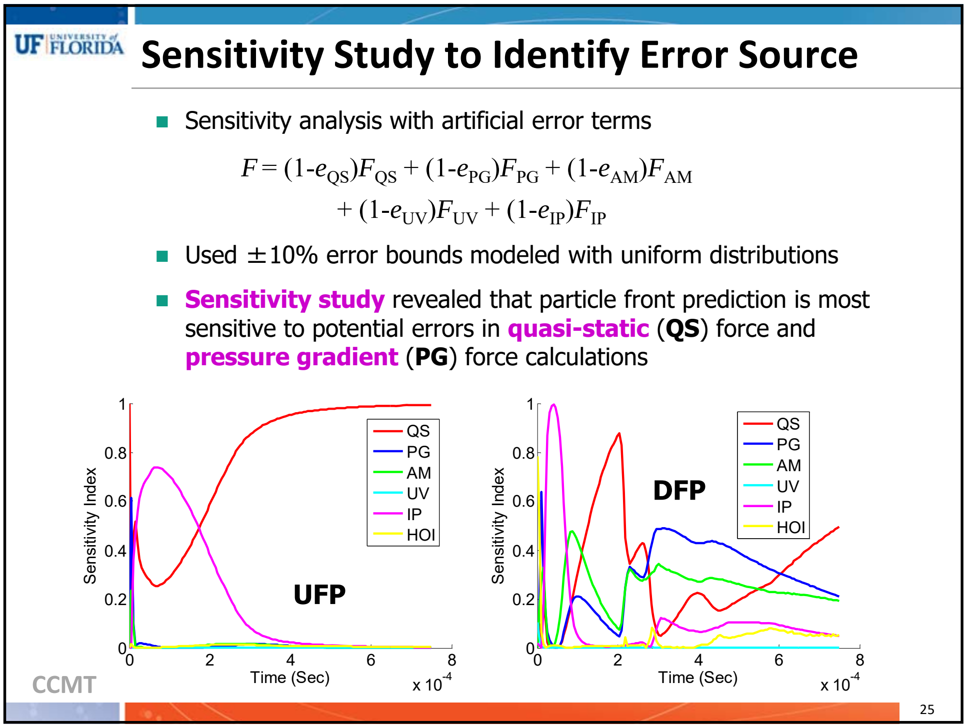

Sensitivity Study to Identify Error Source

- Sensitivity analysis with artificial error terms

$$F = (1-e_{QS})F_{QS} + (1-e_{PG})F_{PG} + (1-e_{AM})F_{AM} + (1-e_{UV})F_{UV} + (1-e_{IP})F_{IP}$$

- Used ±10% error bounds modeled with uniform distributions
- **Sensitivity study** revealed that particle front prediction is most sensitive to potential errors in **quasi-static** (**QS**) force and **pressure gradient** (**PG**) force calculations

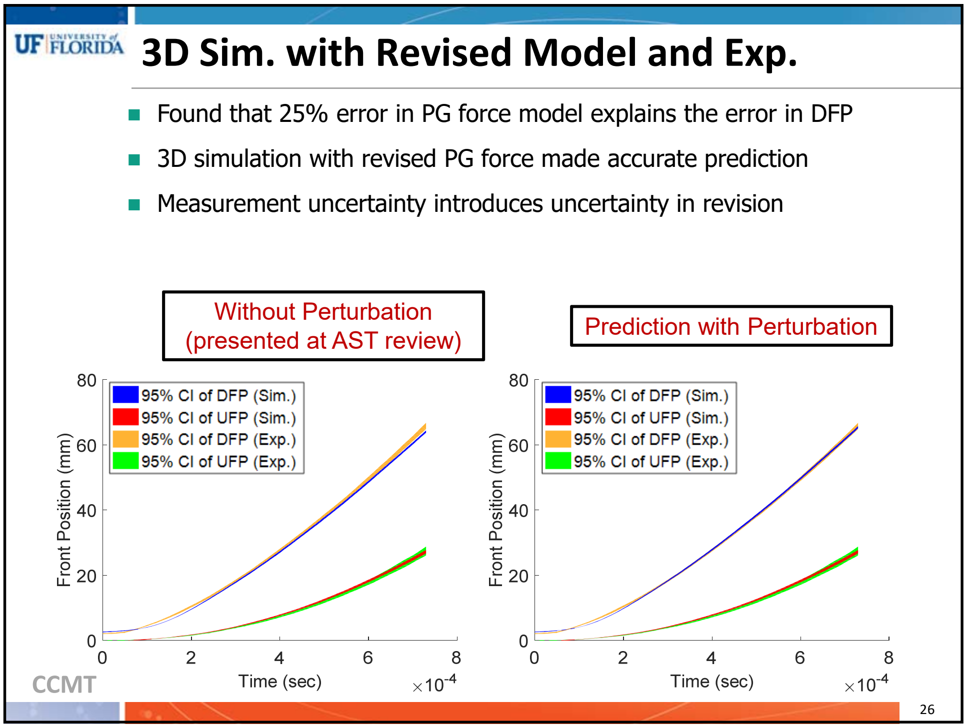

3D Sim. with Revised Model and Exp.

- Found that 25% error in PG force model explains the error in DFP
- 3D simulation with revised PG force made accurate prediction
- Measurement uncertainty introduces uncertainty in revision

Without Perturbation (presented at AST review)

Prediction with Perturbation

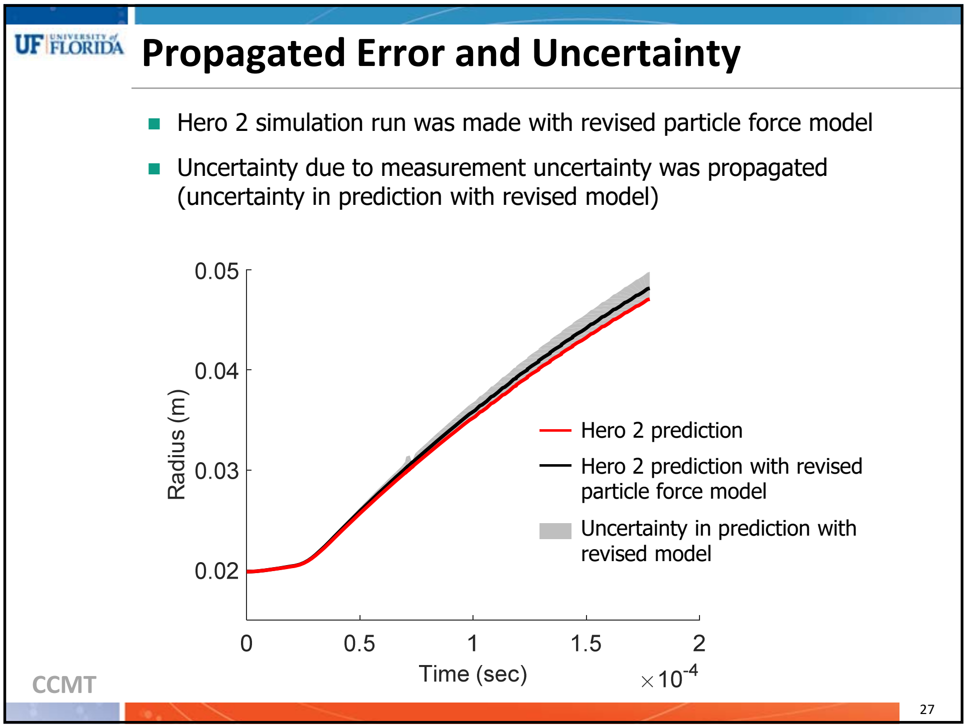

Propagated Error and Uncertainty

- Hero 2 simulation run was made with revised particle force model
- Uncertainty due to measurement uncertainty was propagated (uncertainty in prediction with revised model)

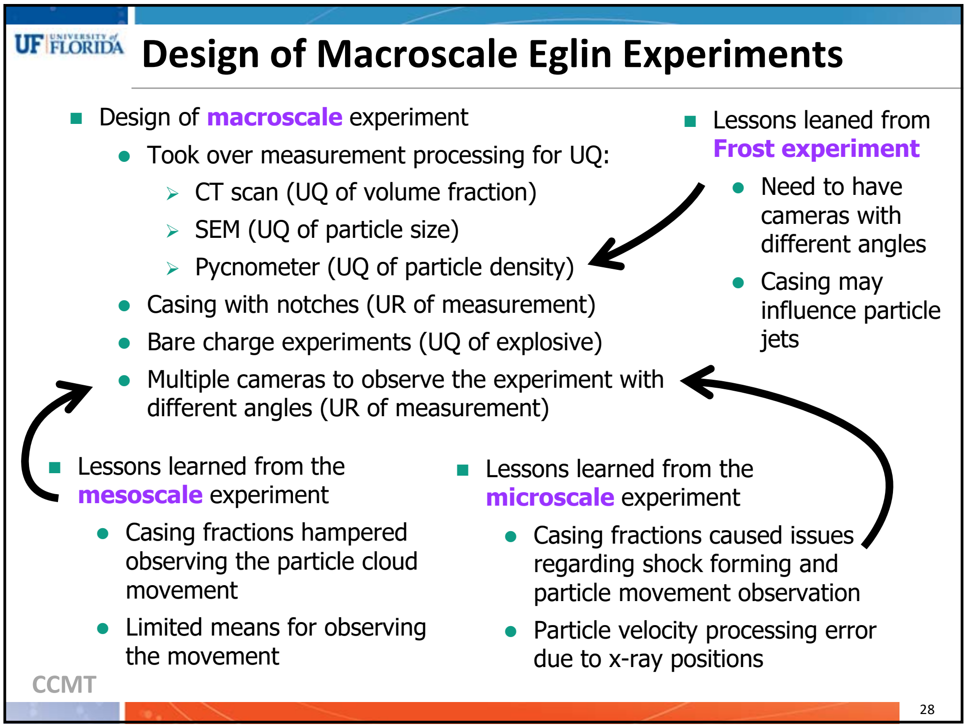

Design of Macroscale Eglin Experiments

- Design of **macroscale** experiment
  - Took over measurement processing for UQ:
    - CT scan (UQ of volume fraction)
    - SEM (UQ of particle size)
    - Pycnometer (UQ of particle density)
  - Casing with notches (UR of measurement)
  - Bare charge experiments (UQ of explosive)
  - Multiple cameras to observe the experiment with different angles (UR of measurement)

- Lessons leaned from **Frost experiment**
  - Need to have cameras with different angles
  - Casing may influence particle jets

- Lessons learned from the **mesoscale** experiment
  - Casing fractions hampered observing the particle cloud movement
  - Limited means for observing the movement

- Lessons learned from the **microscale** experiment
  - Casing fractions caused issues regarding shock forming and particle movement observation
  - Particle velocity processing error due to x-ray positions

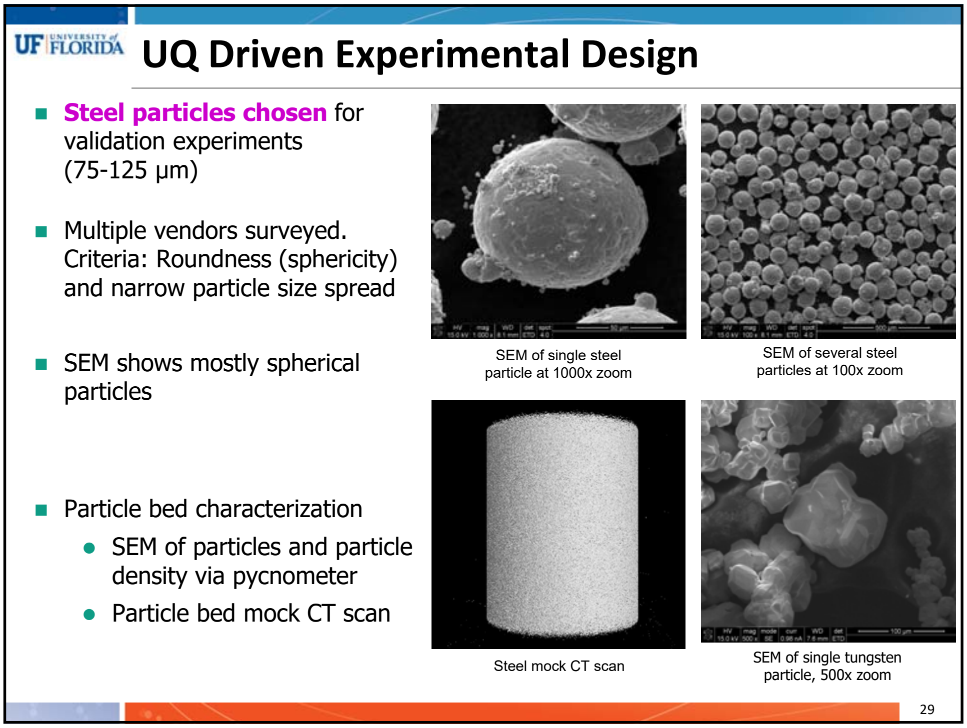

UQ Driven Experimental Design

- **Steel particles chosen** for validation experiments (75-125 µm)
- Multiple vendors surveyed. Criteria: Roundness (sphericity) and narrow particle size spread
- SEM shows mostly spherical particles
- Particle bed characterization
  - SEM of particles and particle density via pycnometer
  - Particle bed mock CT scan

SEM of single steel particle at 1000x zoom

SEM of several steel particles at 100x zoom

Steel mock CT scan

SEM of single tungsten particle, 500x zoom

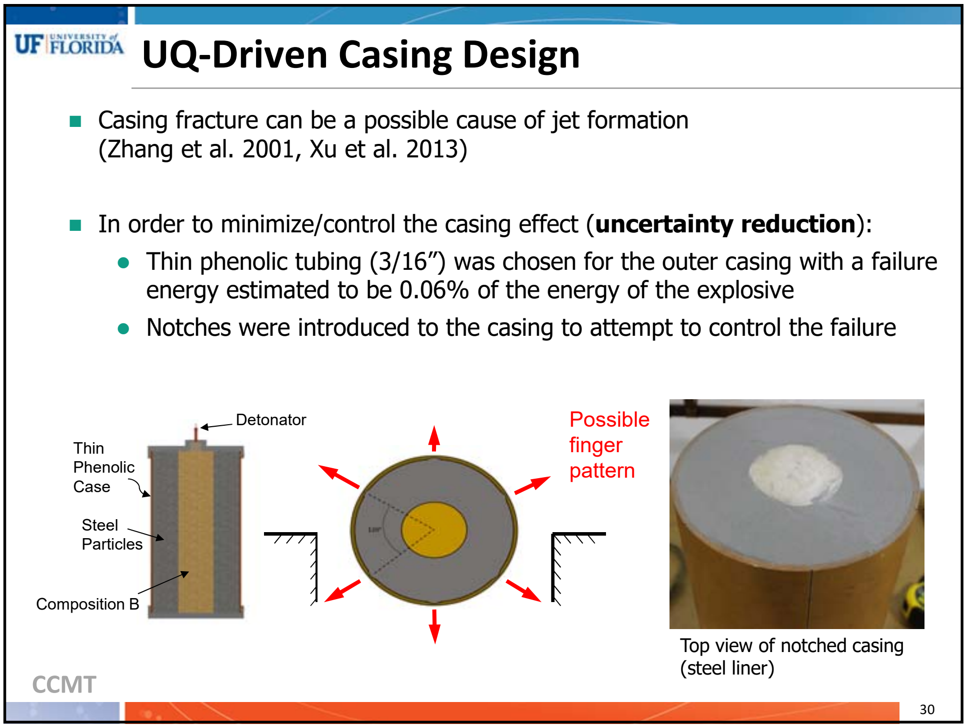

UQ-Driven Casing Design

- Casing fracture can be a possible cause of jet formation (Zhang et al. 2001, Xu et al. 2013)
- In order to minimize/control the casing effect (**uncertainty reduction**):
  - Thin phenolic tubing (3/16”) was chosen for the outer casing with a failure energy estimated to be 0.06% of the energy of the explosive
  - Notches were introduced to the casing to attempt to control the failure

Top view of notched casing (steel liner)

CCMT

Legacy Data and Notching Effect

- Legacy data agrees well with new bare charge results
- Case notches did not appear to have an effect compared to the un-notched data
- Variability is small between 3 data sets (2 tests with notched case and 1 case with un-notched case)
- Jets formed for all cases

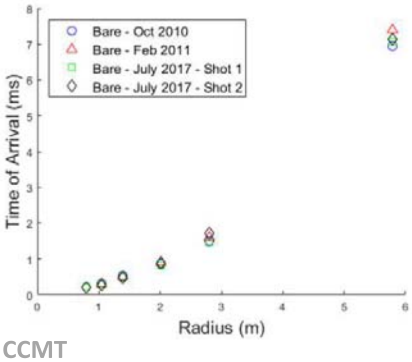

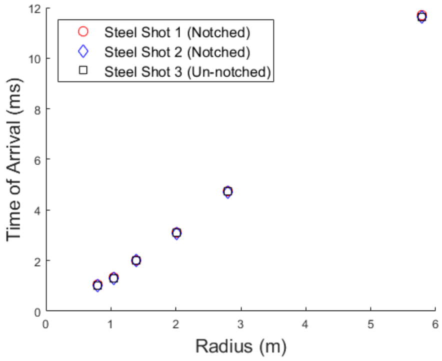

Summary

- Eglin microscale
  - **Calibration** and **UQ** of reactive burn model parameters (knobs) for macroscale
  - **Two most probable calibrated estimates** from different calibration methods with the same data set
  - **Microscale UQ** indicates that uncertainty in the simulation prediction is large and heat release Q is dominant uncertainty source
- UQ of Eglin macroscale
  - **Sensitivity study** identified potential error in pressure gradient force is most likely to be a dominant error source
  - **Propagation** of error and uncertainty in the particle force model calibrates the macroscale prediction and quantifies uncertainty (a lesson learn from lower scale)

CCMT

Do you have any questions?

Backup Slides

Camera 1 vs Camera 3

- Cameras 1 and 4 contain significant perspective errors if used to measure shock position on the centerline
- Camera 3 shows three shock structures due to end-cap effects
- To find the shock time of arrival (TOA) along the 90-degree, camera 3 is used
- Shocks analyzed normal to the ground to examine the shock data for ground effects

Camera 1 at 1.067 ms after detonation

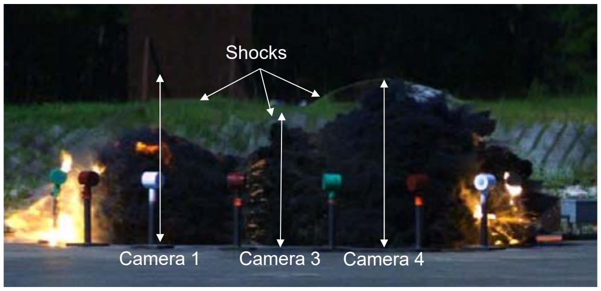

Camera 3 at 1.067 ms after detonation

Transverse Shock Positions

- Transverse shock positions for RZ thickness 1.7 mm and 0.88 mm
- RZ thickness 0.88 predicts higher shock position but the **initial positions are different**

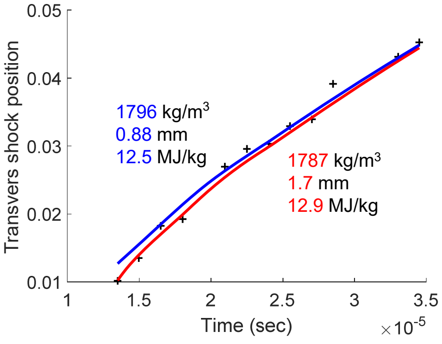

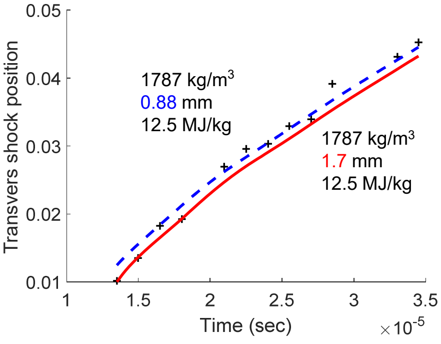

CCMT

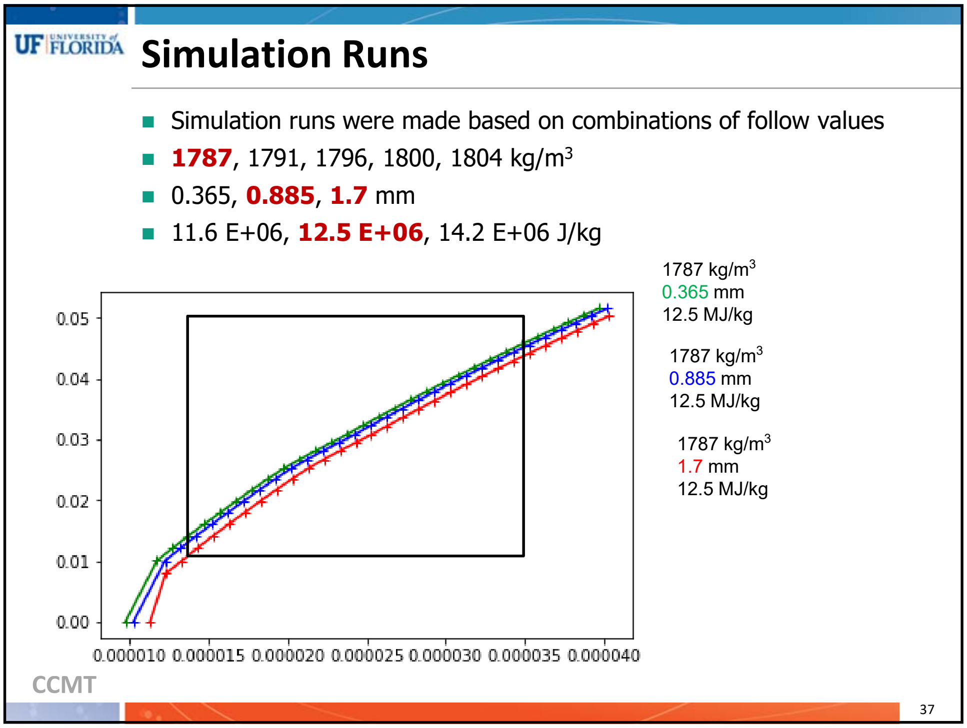

Simulation Runs

- Simulation runs were made based on combinations of follow values
- **1787**, 1791, 1796, 1800, 1804 kg/m^3
- 0.365, **0.885**, **1.7** mm
- 11.6 E+06, **12.5 E+06**, 14.2 E+06 J/kg

1787 kg/m^3
0.365 mm
12.5 MJ/kg

1787 kg/m^3
0.885 mm
12.5 MJ/kg

1787 kg/m^3
1.7 mm
12.5 MJ/kg

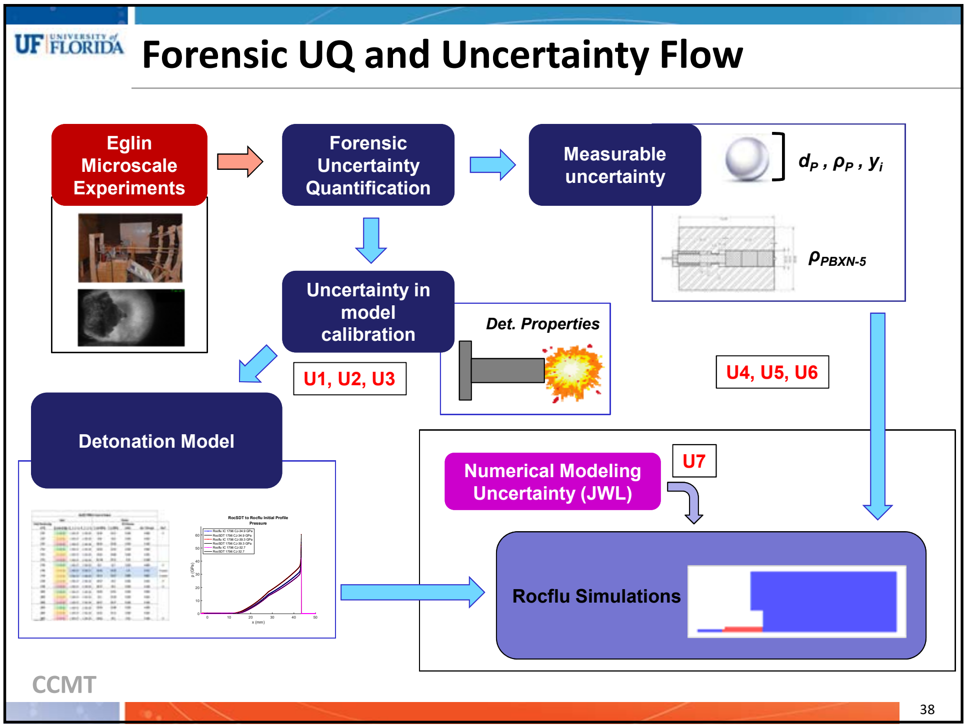

Forensic UQ and Uncertainty Flow

Eglin Microscale Experiments

Forensic Uncertainty Quantification

Measurable uncertainty

d_P , ρ_P , y_i

$\rho_{PBXN\text{-}5}$

Uncertainty in model calibration

U1, U2, U3

Det. Properties

U4, U5, U6

Detonation Model

Numerical Modeling Uncertainty (JWL)

U7

Rocflu Simulations

Uncertainties

Experimental Uncertainties

- Particle diameter: mean = 2.016 mm, st. dev = 0.008 mm
- Particle density: mean = 15.8 g/cc, st. dev = 0.07 g/cc

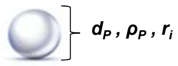

- The initial radial position of the particle was varied to account for machining tolerance of the explosive assembly (± 0.254 mm)
- Explosive density (PBXN-5): mean = 1796 kg/m^3, st. dev = 3 kg/m^3

Reactive Burn Model Parameter Uncertainty

- Initial conditions for simulation of the experiment consider the explosive at the moment of complete detonation
- An external code is used to simulate the detonation of the explosive pellets
- Modeling the explosive results in additional uncertainties

Explosive Modeling Uncertainty

Inputs for detonation simulations.

Density (PBXN-5)	[1787,1804] kg/m^3
Heat of Reaction	[11.6, 14.2] MJ/kg
Reaction Zone Thickness	[0.365,1.700] mm

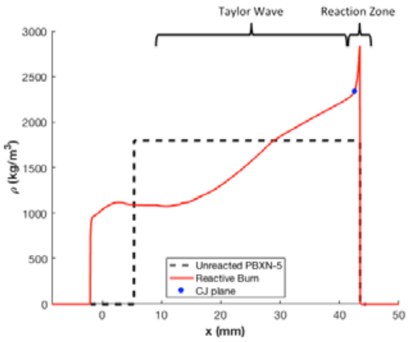

Explosive Modeling Uncertainty

- Detonation simulation uncertain inputs:
  - Explosive density, obtained from laboratory measurements
  - Heat of reaction was calibrated for the pressure solution to reach the CJ pressure for PBXN-5 in the literature (32.7 to 39.3 GPa)
  - Reaction rate parameters govern the reaction zone thickness, the literature gives this thickness as 0.1 to 0.7 mm for PBXN-5

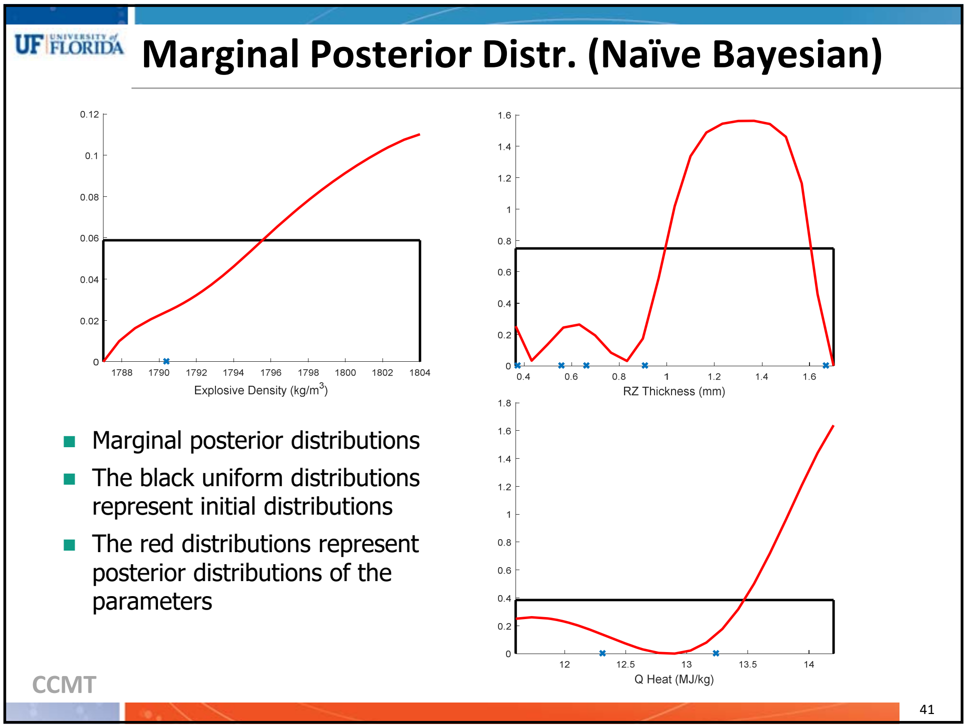

Marginal Posterior Distr. (Naïve Bayesian)

- Marginal posterior distributions
- The black uniform distributions represent initial distributions
- The red distributions represent posterior distributions of the parameters

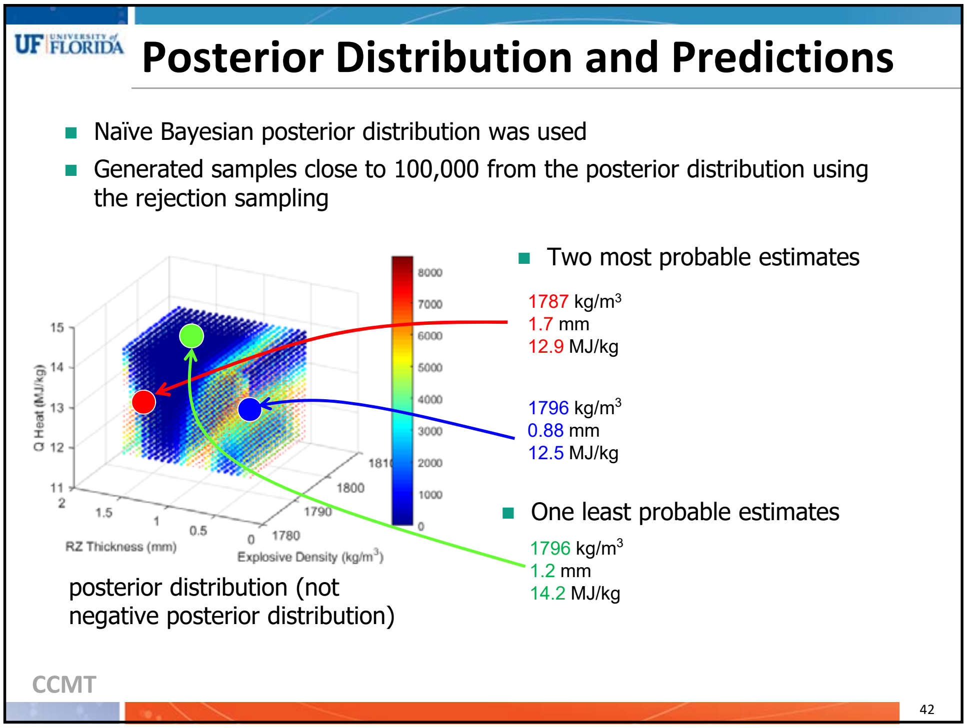

Posterior Distribution and Predictions

- Naïve Bayesian posterior distribution was used
- Generated samples close to 100,000 from the posterior distribution using the rejection sampling

posterior distribution (not negative posterior distribution)

- Two most probable estimates

1787 kg/m^3
1.7 mm
12.9 MJ/kg

1796 kg/m^3
0.88 mm
12.5 MJ/kg

- One least probable estimates

1796 kg/m^3
1.2 mm
14.2 MJ/kg

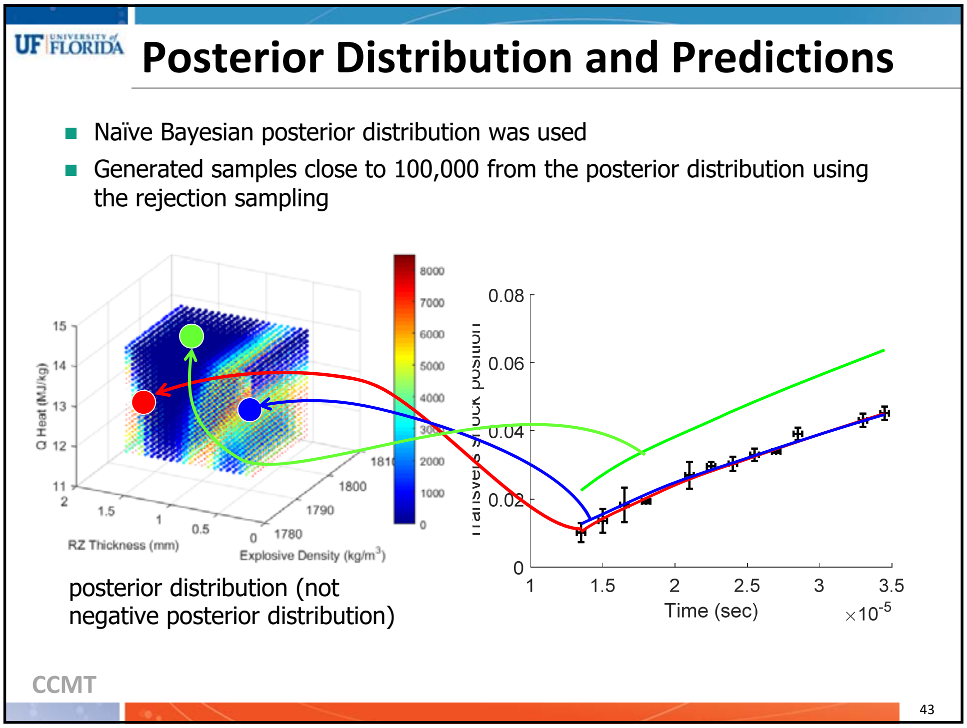

Posterior Distribution and Predictions

- Naïve Bayesian posterior distribution was used
- Generated samples close to 100,000 from the posterior distribution using the rejection sampling

posterior distribution (not negative posterior distribution)

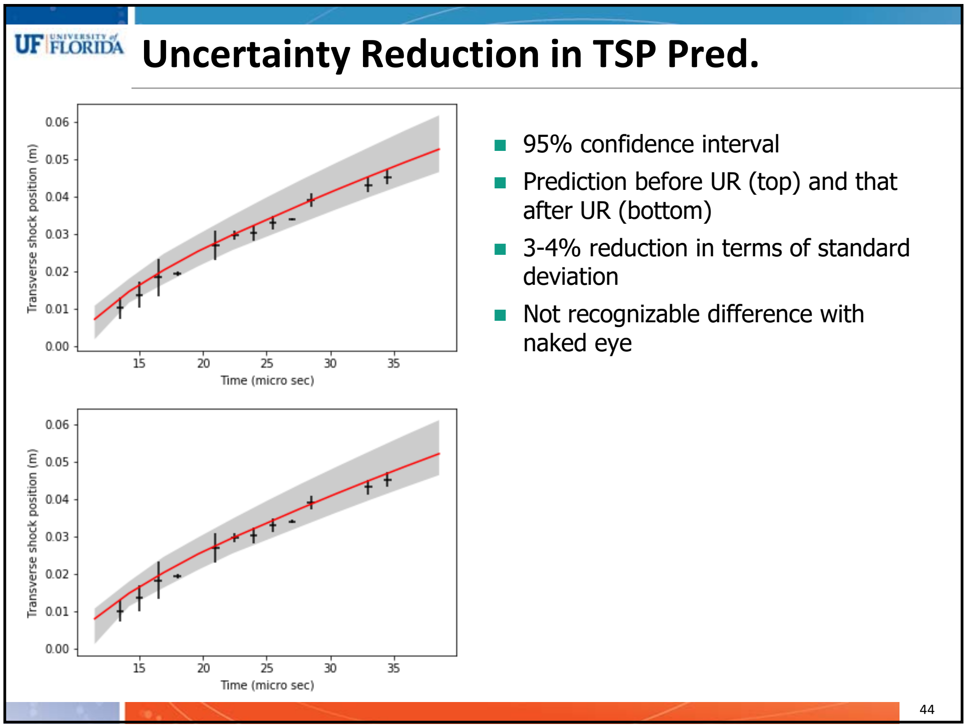

Uncertainty Reduction in TSP Pred.

- 95% confidence interval
- Prediction before UR (top) and that after UR (bottom)
- 3-4% reduction in terms of standard deviation
- Not recognizable difference with naked eye

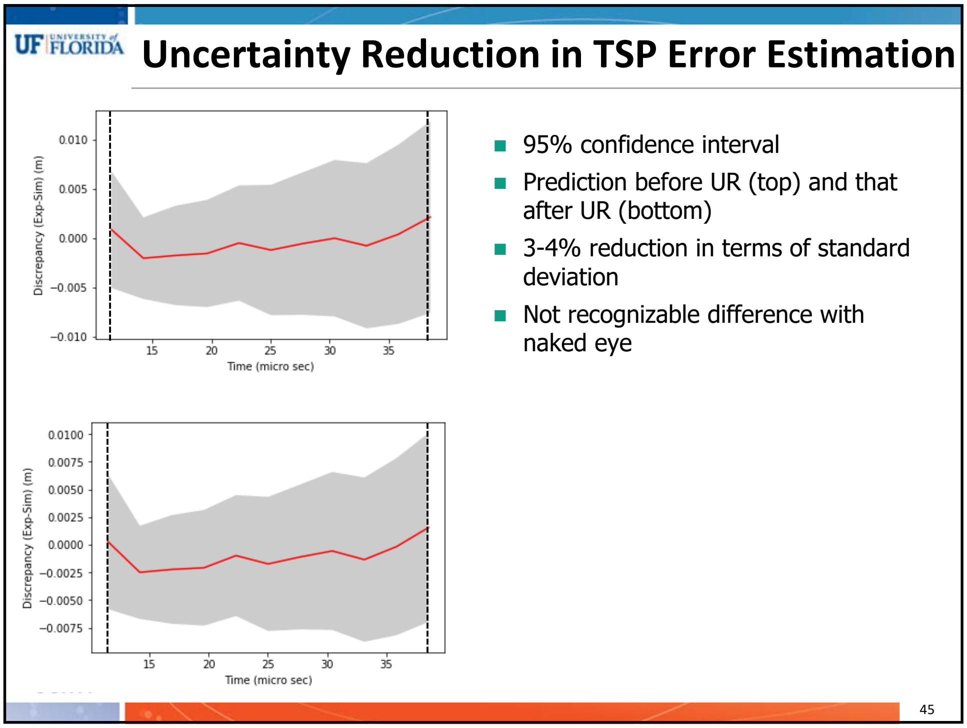

Uncertainty Reduction in TSP Error Estimation

- 95% confidence interval
- Prediction before UR (top) and that after UR (bottom)
- 3-4% reduction in terms of standard deviation
- Not recognizable difference with naked eye

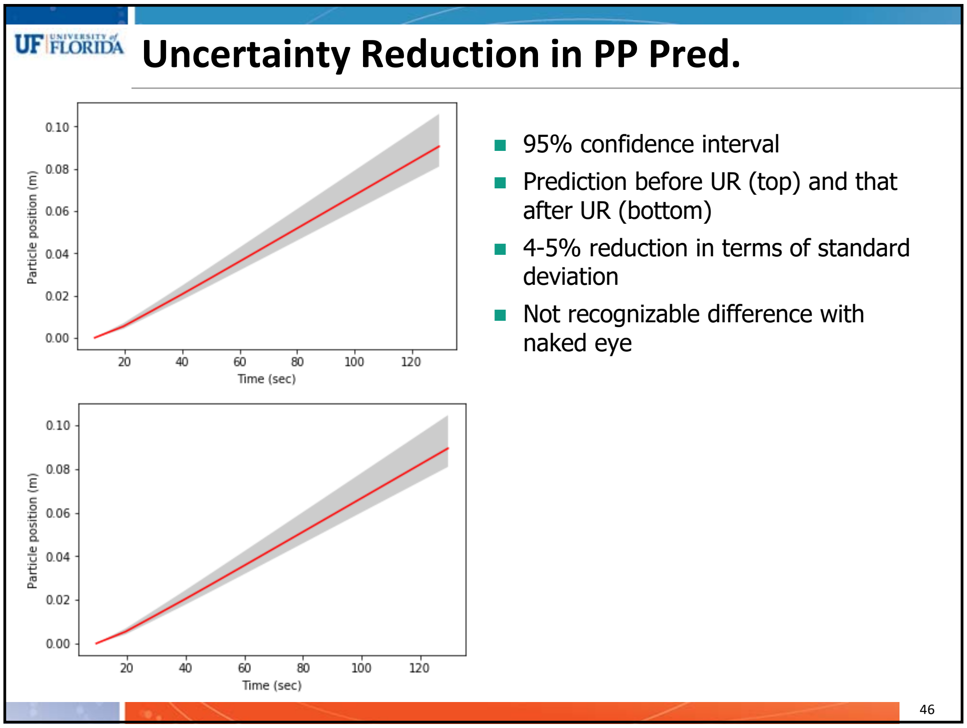

Uncertainty Reduction in PP Pred.

- 95% confidence interval
- Prediction before UR (top) and that after UR (bottom)
- 4-5% reduction in terms of standard deviation
- Not recognizable difference with naked eye

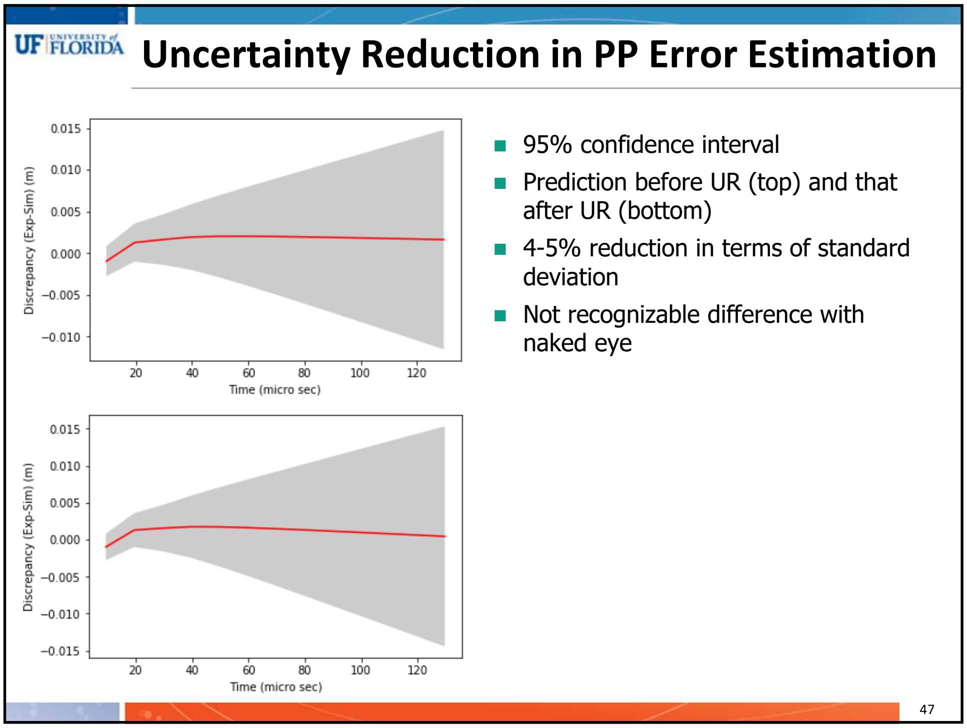

Uncertainty Reduction in PP Error Estimation

- 95% confidence interval
- Prediction before UR (top) and that after UR (bottom)
- 4-5% reduction in terms of standard deviation
- Not recognizable difference with naked eye

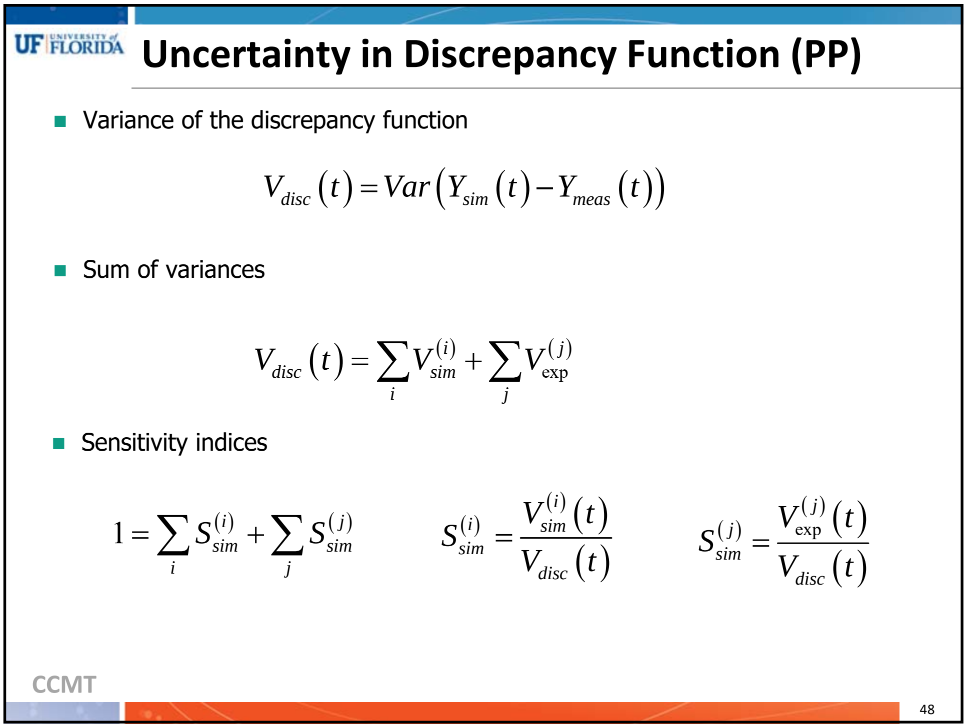

Uncertainty in Discrepancy Function (PP)

- Variance of the discrepancy function

$$V_{disc}\left(t\right)=Var\left(Y_{sim}\left(t\right)-Y_{meas}\left(t\right)\right)$$

- Sum of variances

$$V_{disc}\left(t\right)=\sum_{i}V_{\text{sim}}^{(i)}+\sum_{j}V_{\text{exp}}^{(j)}$$

- Sensitivity indices

$$1=\sum_{i}S_{\text{sim}}^{(i)}+\sum_{j}S_{\text{sim}}^{(j)} \qquad S_{\text{sim}}^{(i)}=\frac{V_{\text{sim}}^{(i)}\left(t\right)}{V_{disc}\left(t\right)} \qquad S_{\text{sim}}^{(j)}=\frac{V_{\text{exp}}^{(j)}\left(t\right)}{V_{disc}\left(t\right)}$$

CCMT

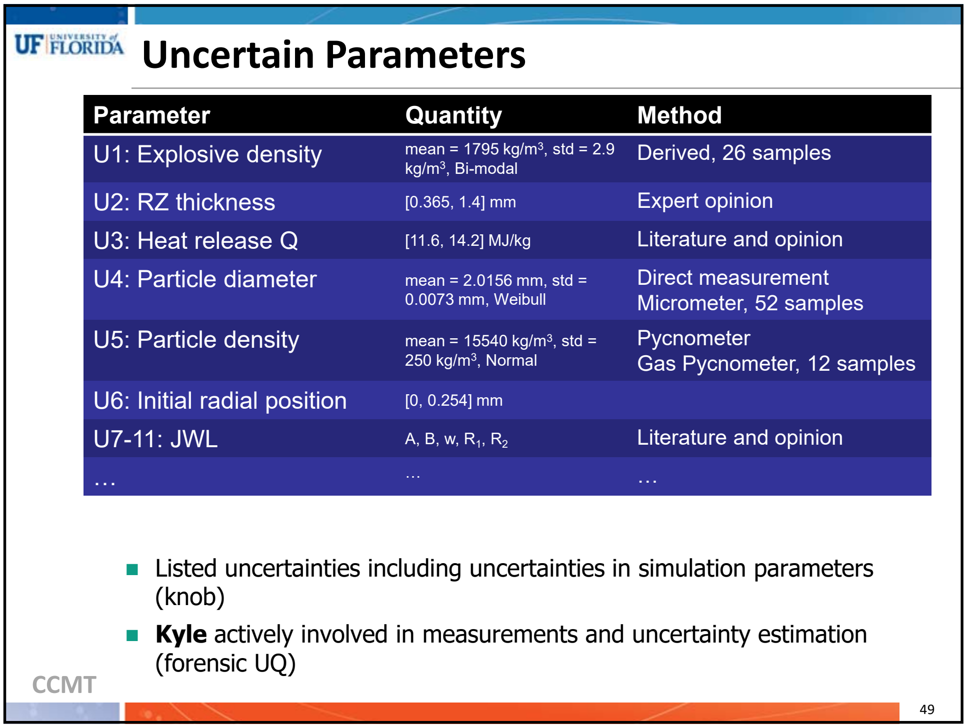

Uncertain Parameters

Parameter	Quantity	Method
U1: Explosive density	mean = 1795 kg/m^3, std = 2.9 kg/m^3, Bi-modal	Derived, 26 samples
U2: RZ thickness	[0.365, 1.4] mm	Expert opinion
U3: Heat release Q	[11.6, 14.2] MJ/kg	Literature and opinion
U4: Particle diameter	mean = 2.0156 mm, std = 0.0073 mm, Weibull	Direct measurement Micrometer, 52 samples
U5: Particle density	mean = 15540 kg/m^3, std = 250 kg/m^3, Normal	Pycnometer Gas Pycnometer, 12 samples
U6: Initial radial position	[0, 0.254] mm	
U7-11: JWL	A, B, w, R$_1$, R$_2$	Literature and opinion
…	…	…

- Listed uncertainties including uncertainties in simulation parameters (knob)
- **Kyle** actively involved in measurements and uncertainty estimation (forensic UQ)

CCMT

CMT-nek: CS Update

Tania Banerjee
Computer and Information Science and Engineering
University of Florida
tmishra@ufl.edu

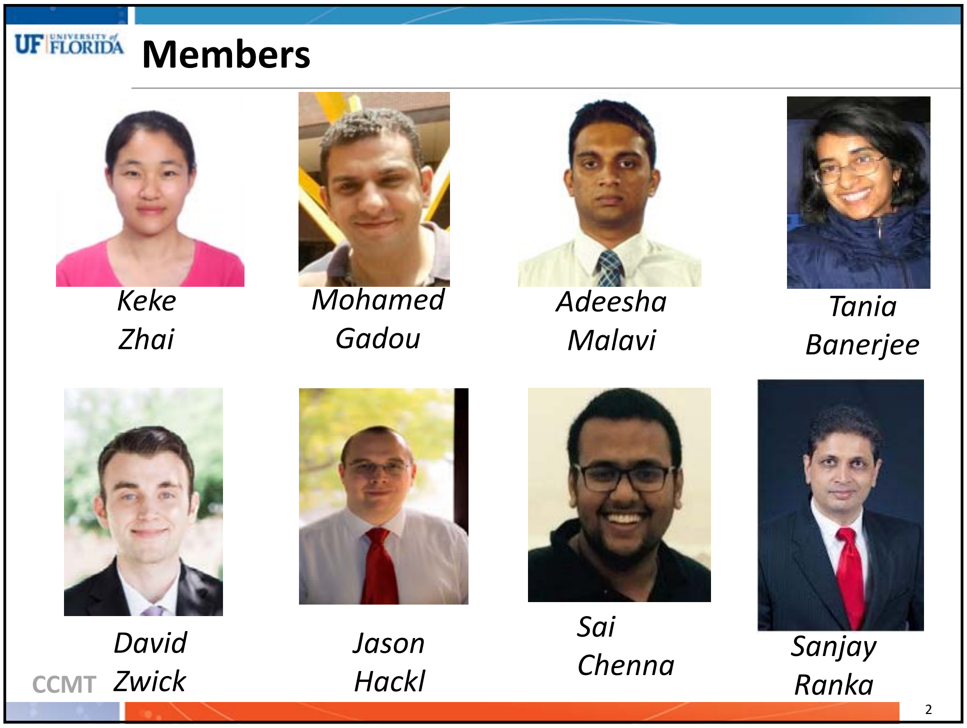

Members

Keke Zhai

Mohamed Gadou

Adeesha Malavi

Tania Banerjee

David Zwick

Jason Hackl

Sai Chenna

Sanjay Ranka

CCMT

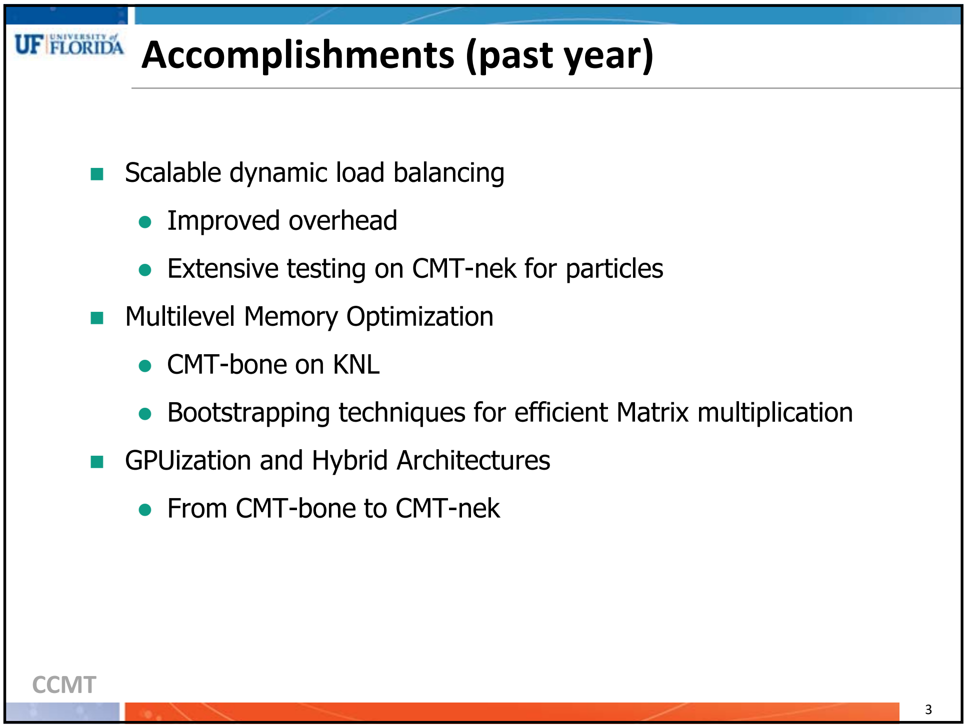

Accomplishments (past year)

- Scalable dynamic load balancing
  - Improved overhead
  - Extensive testing on CMT-nek for particles
- Multilevel Memory Optimization
  - CMT-bone on KNL
  - Bootstrapping techniques for efficient Matrix multiplication
- GPUization and Hybrid Architectures
  - From CMT-bone to CMT-nek

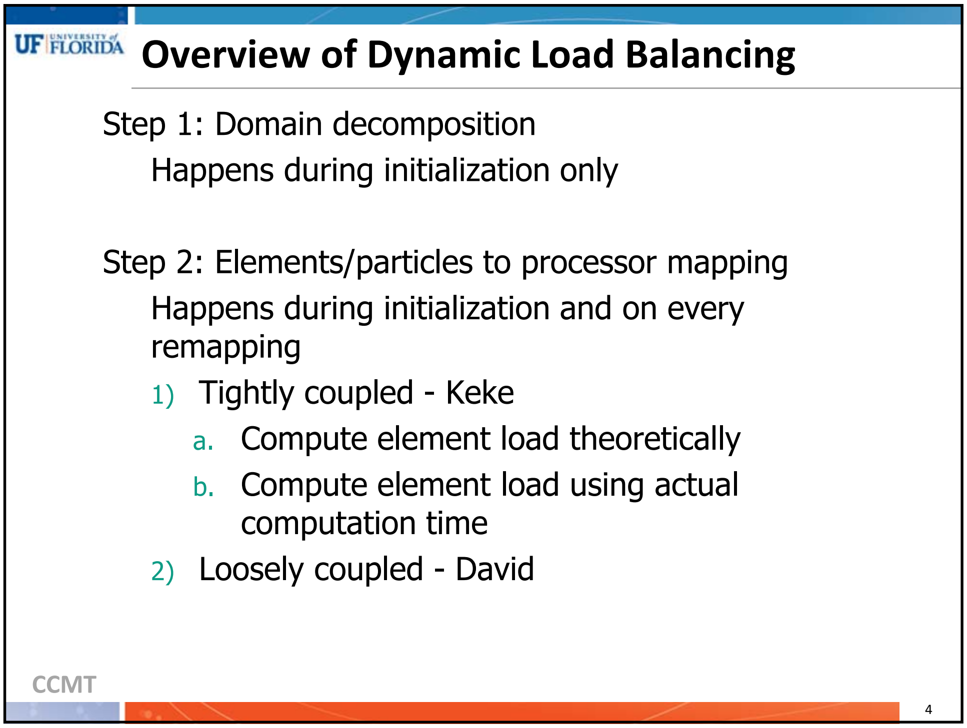

Overview of Dynamic Load Balancing

Step 1: Domain decomposition

Happens during initialization only

Step 2: Elements/particles to processor mapping

Happens during initialization and on every remapping

1) Tightly coupled - Keke
   a. Compute element load theoretically
   b. Compute element load using actual computation time
2) Loosely coupled - David

Overview of Dynamic Load Balancing

Step 3: Decide when to trigger a remap

1. Rebalance after every k time steps (user set up)
2. Rebalance automatically after certain time steps (adaptive load balancing)
3. Rebalance when the difference in time between the fastest and the slowest processor exceeds a threshold.

Step 4: Transfer elements and particles and reset other data structures

---

Domain Decomposition

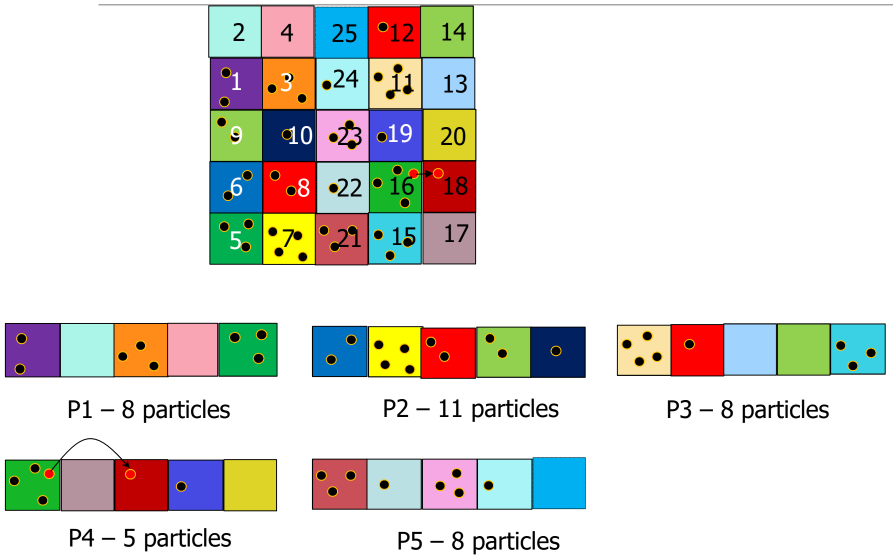

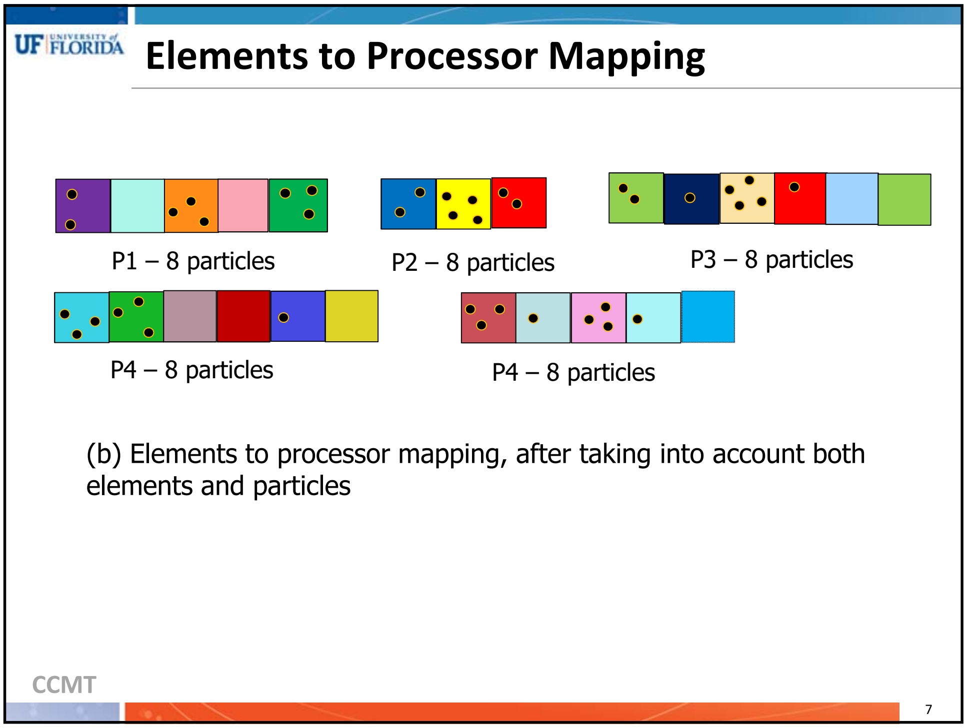

Elements to Processor Mapping

P1 – 8 particles

P2 – 8 particles

P3 – 8 particles

P4 – 8 particles

P4 – 8 particles

(b) Elements to processor mapping, after taking into account both elements and particles

CCMT

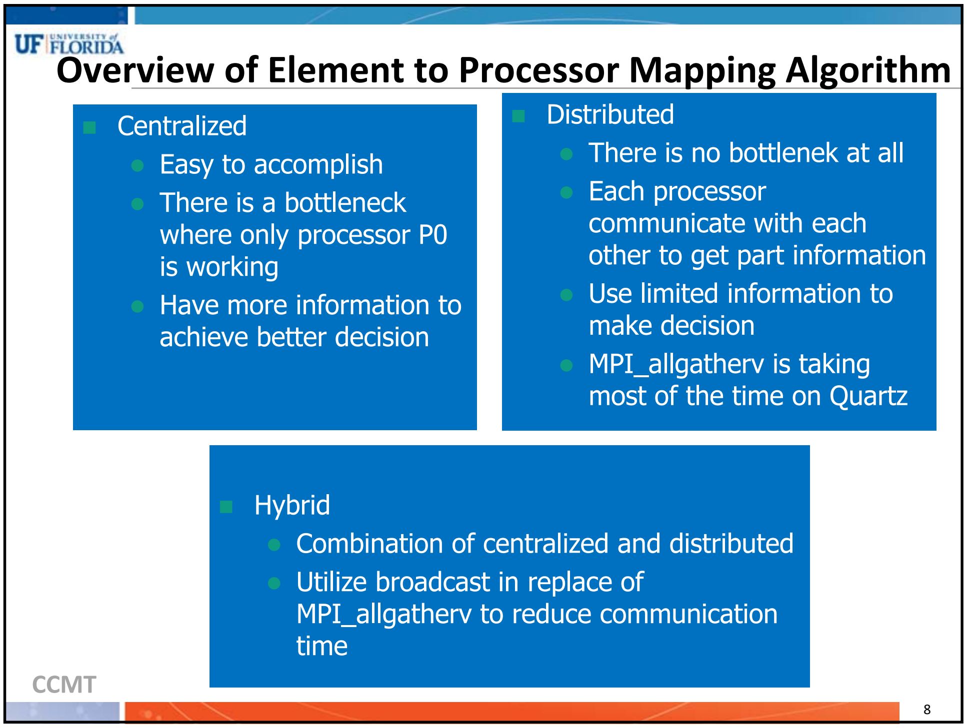

Overview of Element to Processor Mapping Algorithm

- Centralized
  - Easy to accomplish
  - There is a bottleneck where only processor P0 is working
  - Have more information to achieve better decision

- Distributed
  - There is no bottlenek at all
  - Each processor communicate with each other to get part information
  - Use limited information to make decision
  - MPI_allgatherv is taking most of the time on Quartz

- Hybrid
  - Combination of centralized and distributed
  - Utilize broadcast in replace of MPI_allgatherv to reduce communication time

CCMT

The Centralized Algorithm

Each processor sends element load array to P0

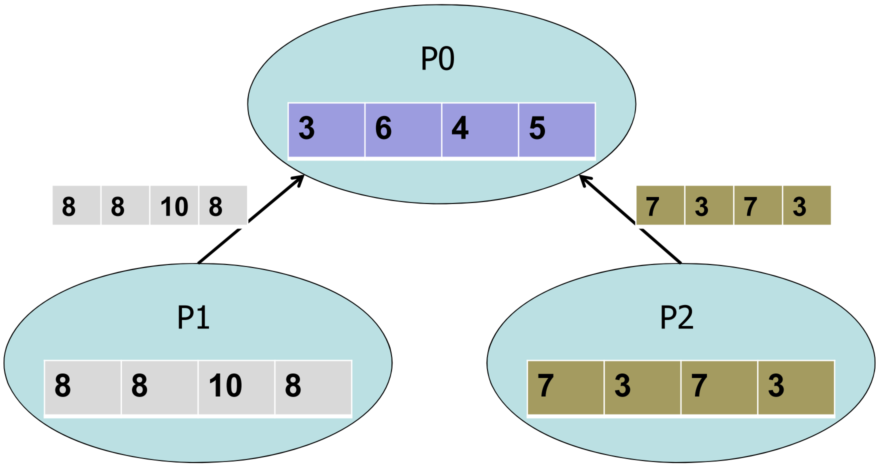

The Distributed Algorithm

Each processor computes local prefix sum and the exclusive prefix sum of the element load on each processor

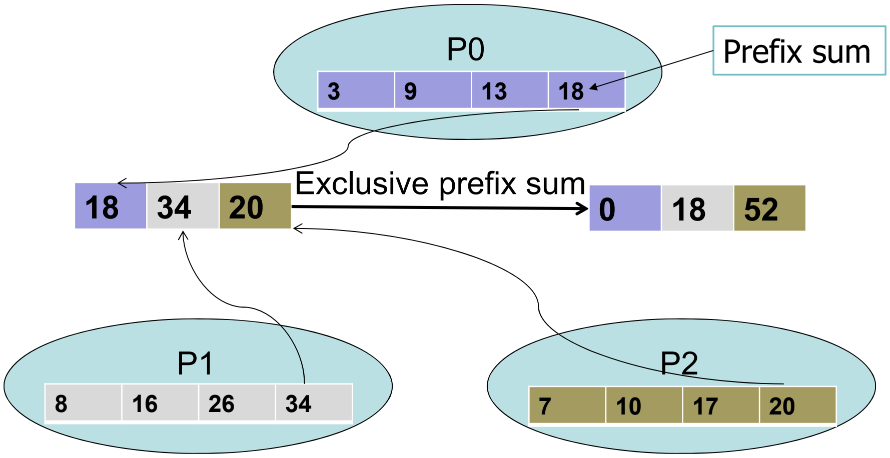

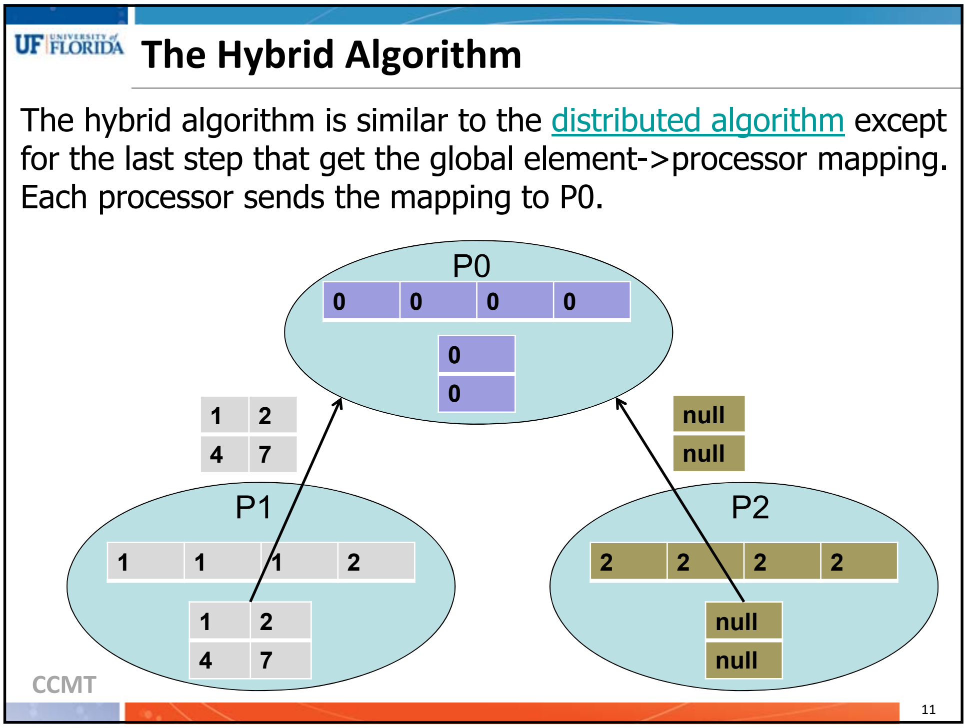

The Hybrid Algorithm

The hybrid algorithm is similar to the distributed algorithm except for the last step that get the global element->processor mapping. Each processor sends the mapping to P0.

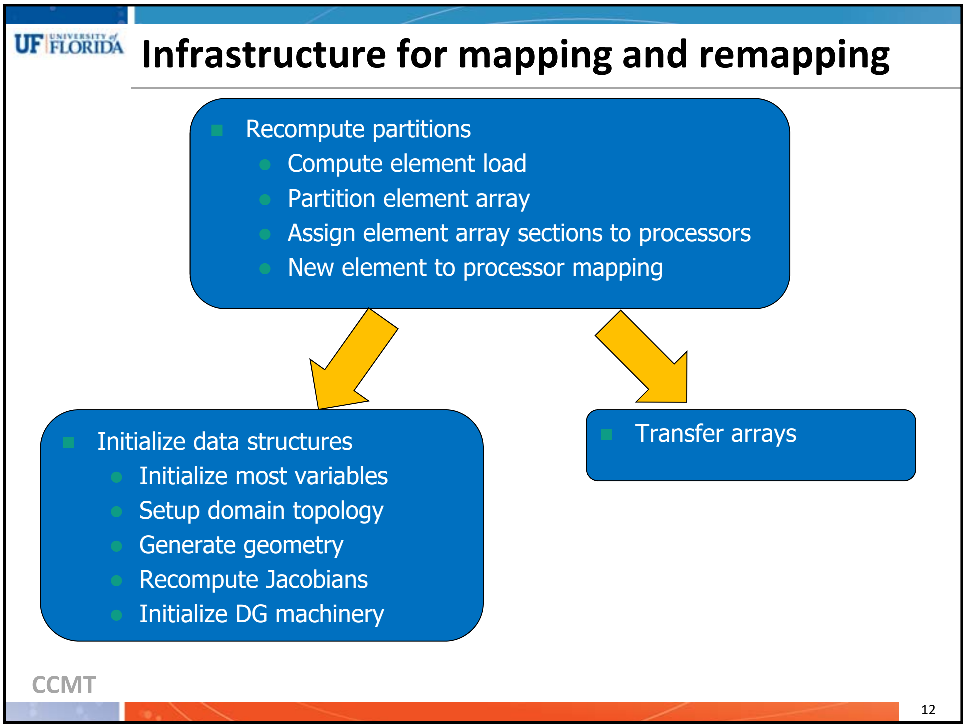

Infrastructure for mapping and remapping

- Recompute partitions
  - Compute element load
  - Partition element array
  - Assign element array sections to processors
  - New element to processor mapping

- Initialize data structures
  - Initialize most variables
  - Setup domain topology
  - Generate geometry
  - Recompute Jacobians
  - Initialize DG machinery

- Transfer arrays

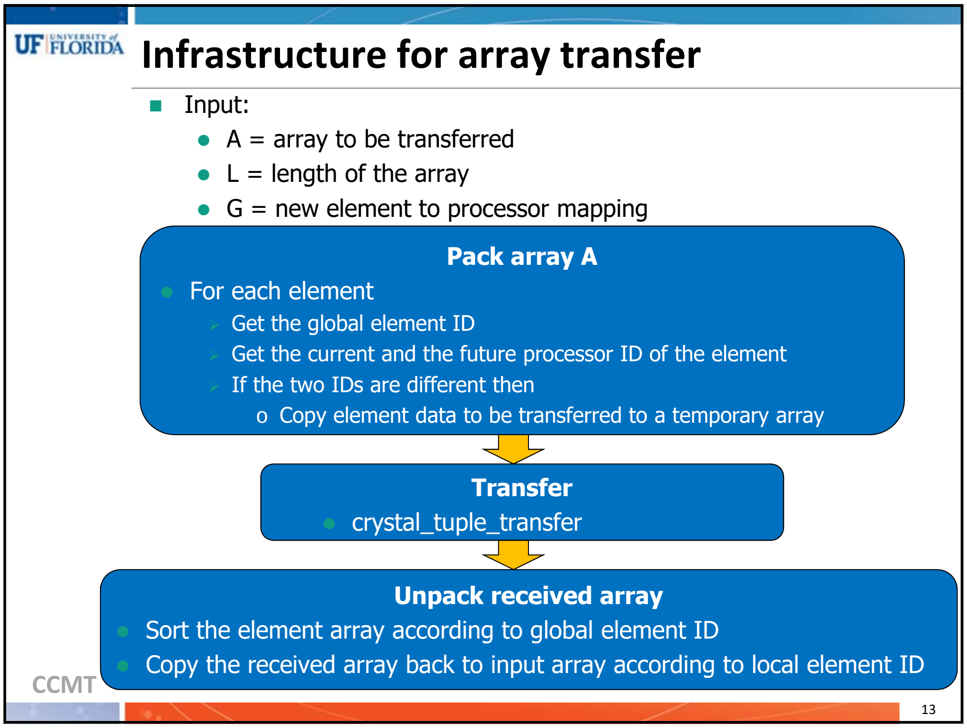

Infrastructure for array transfer

- Input:
  - A = array to be transferred
  - L = length of the array
  - G = new element to processor mapping

Pack array A

- For each element
  - Get the global element ID
  - Get the current and the future processor ID of the element
  - If the two IDs are different then
    - Copy element data to be transferred to a temporary array

Transfer

- crystal_tuple_transfer

Unpack received array

- Sort the element array according to global element ID
- Copy the received array back to input array according to local element ID

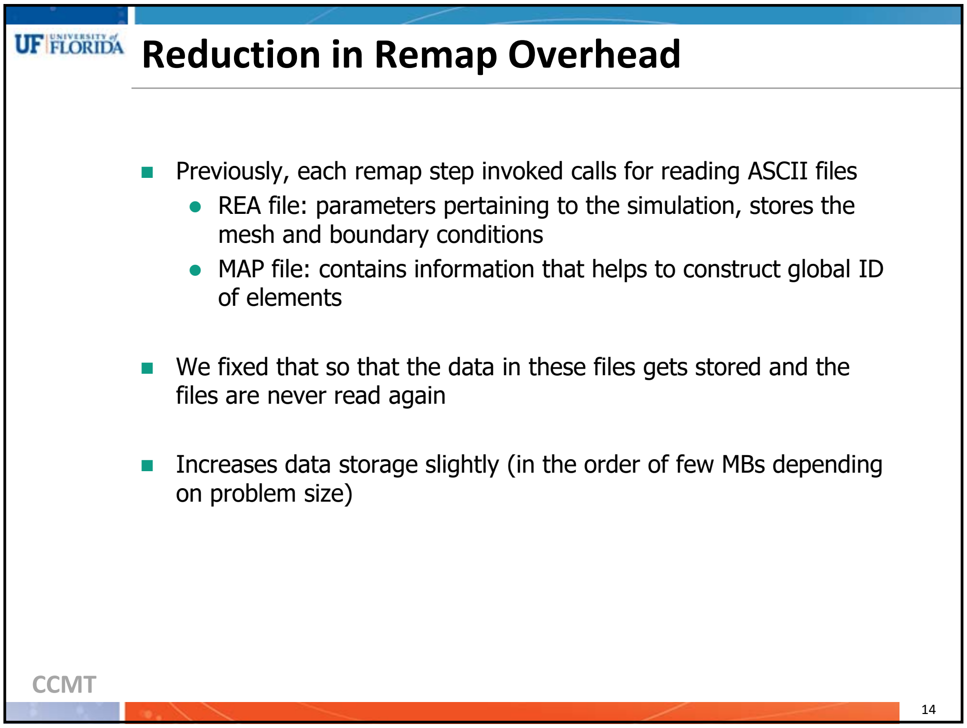

Reduction in Remap Overhead

- Previously, each remap step invoked calls for reading ASCII files
  - REA file: parameters pertaining to the simulation, stores the mesh and boundary conditions
  - MAP file: contains information that helps to construct global ID of elements
- We fixed that so that the data in these files gets stored and the files are never read again
- Increases data storage slightly (in the order of few MBs depending on problem size)

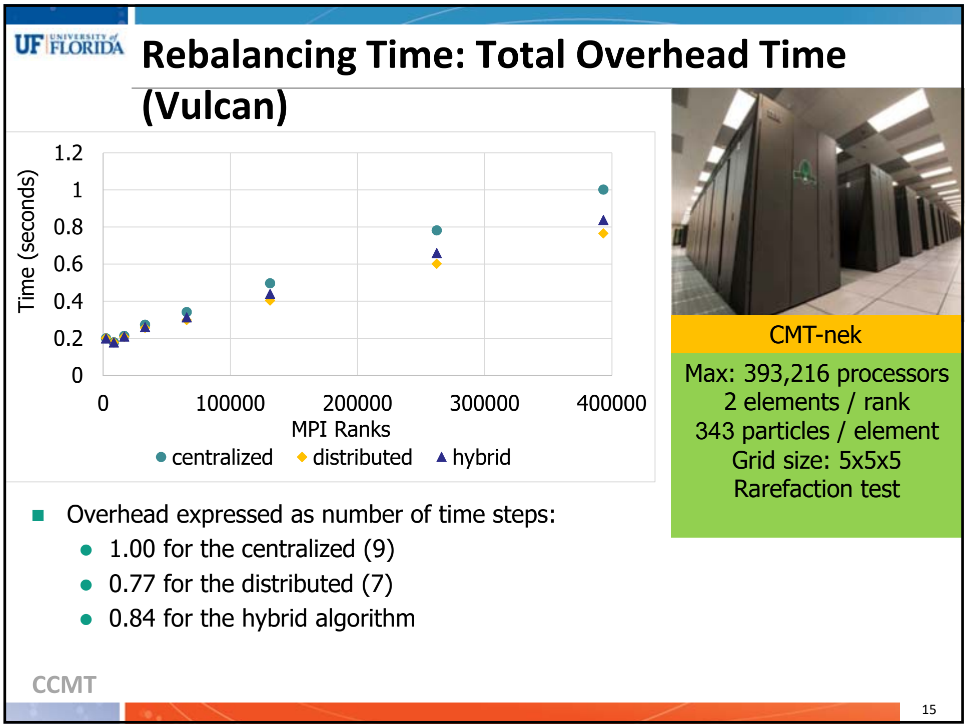

Rebalancing Time: Total Overhead Time (Vulcan)

CMT-nek

Max: 393,216 processors
2 elements / rank
343 particles / element
Grid size: 5x5x5
Rarefaction test

- Overhead expressed as number of time steps:
  - 1.00 for the centralized (9)
  - 0.77 for the distributed (7)
  - 0.84 for the hybrid algorithm

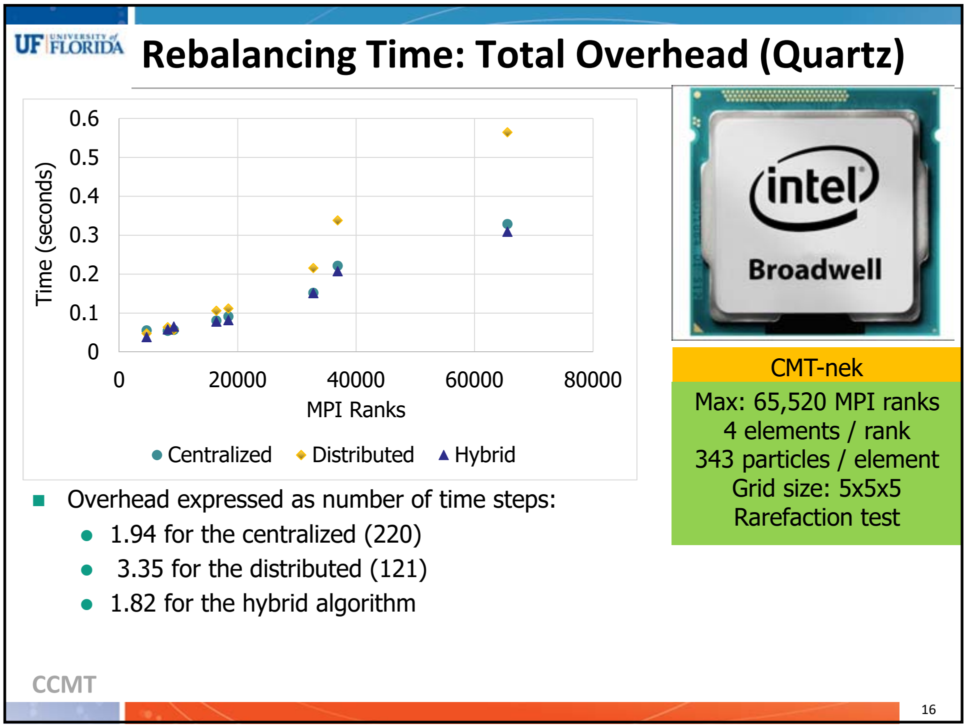

Rebalancing Time: Total Overhead (Quartz)

CMT-nek

Max: 65,520 MPI ranks
4 elements / rank
343 particles / element
Grid size: 5x5x5
Rarefaction test

- Overhead expressed as number of time steps:
  - 1.94 for the centralized (220)
  - 3.35 for the distributed (121)
  - 1.82 for the hybrid algorithm

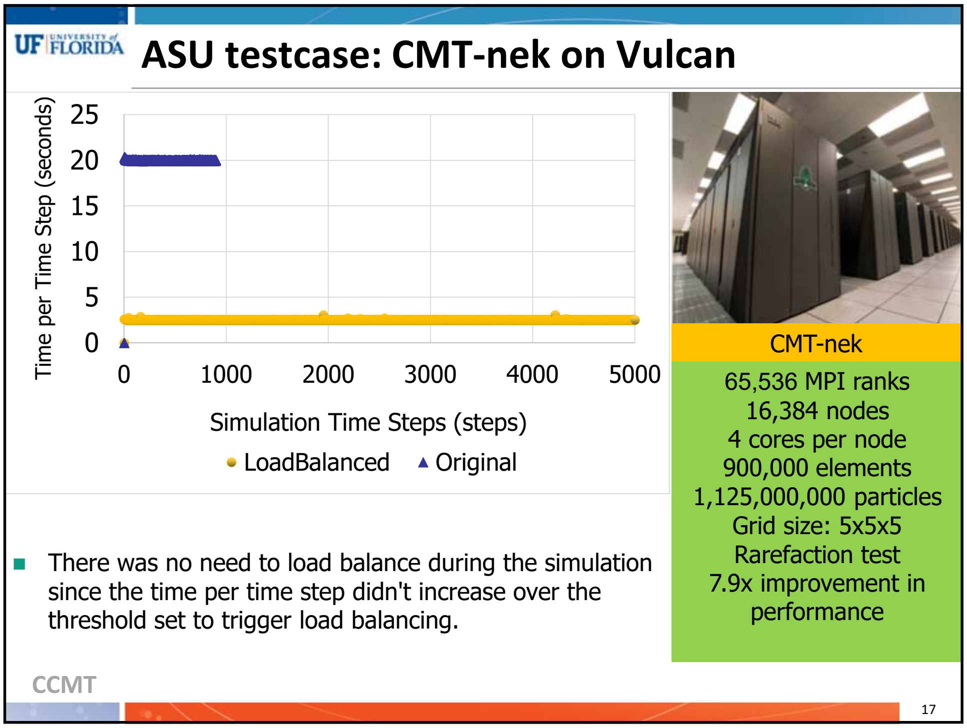

ASU testcase: CMT-nek on Vulcan

CMT-nek

65,536 MPI ranks
16,384 nodes
4 cores per node
900,000 elements
1,125,000,000 particles
Grid size: 5x5x5
Rarefaction test
7.9x improvement in performance

- There was no need to load balance during the simulation since the time per time step didn't increase over the threshold set to trigger load balancing.

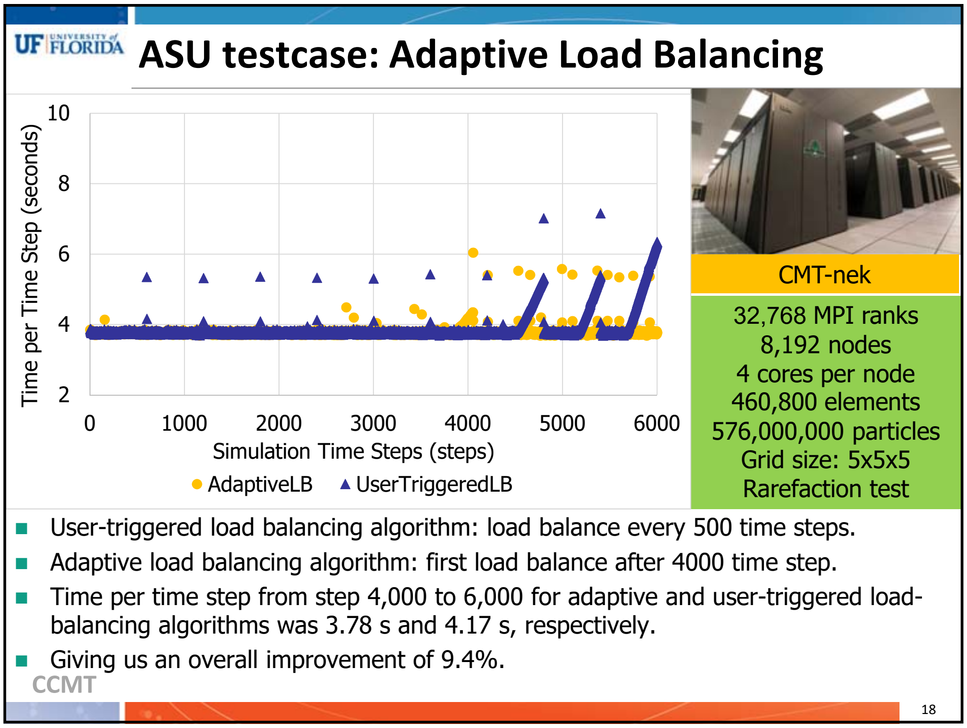

ASU testcase: Adaptive Load Balancing

CMT-nek

32,768 MPI ranks
8,192 nodes
4 cores per node
460,800 elements
576,000,000 particles
Grid size: 5x5x5
Rarefaction test

- User-triggered load balancing algorithm: load balance every 500 time steps.
- Adaptive load balancing algorithm: first load balance after 4000 time step.
- Time per time step from step 4,000 to 6,000 for adaptive and user-triggered load-balancing algorithms was 3.78 s and 4.17 s, respectively.
- Giving us an overall improvement of 9.4%.

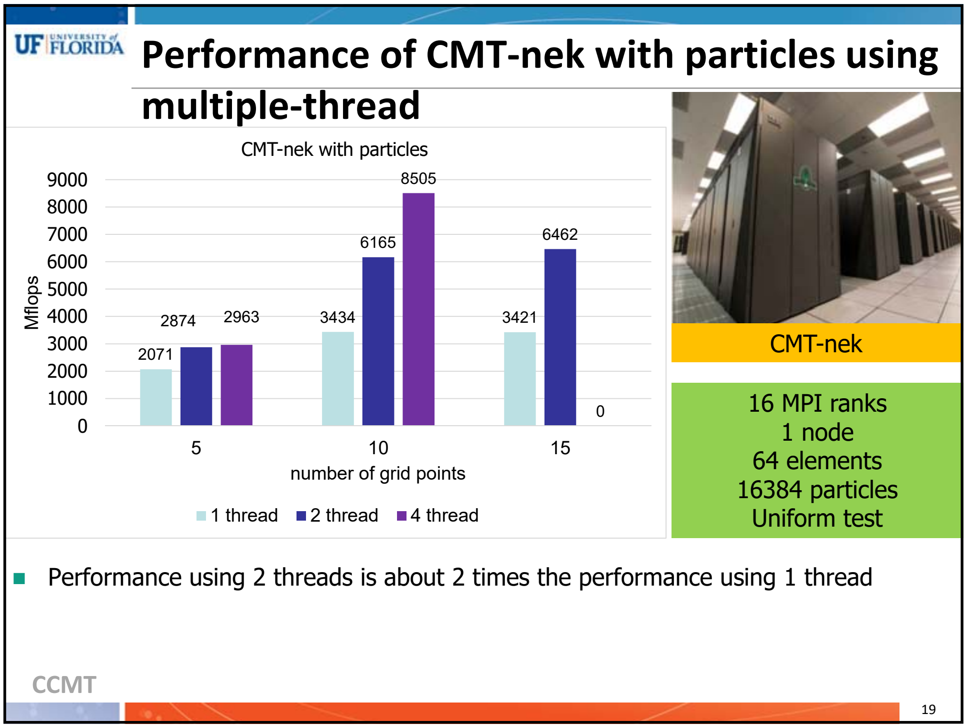

Performance of CMT-nek with particles using multiple-thread

16 MPI ranks
1 node
64 elements
16384 particles
Uniform test

- Performance using 2 threads is about 2 times the performance using 1 thread

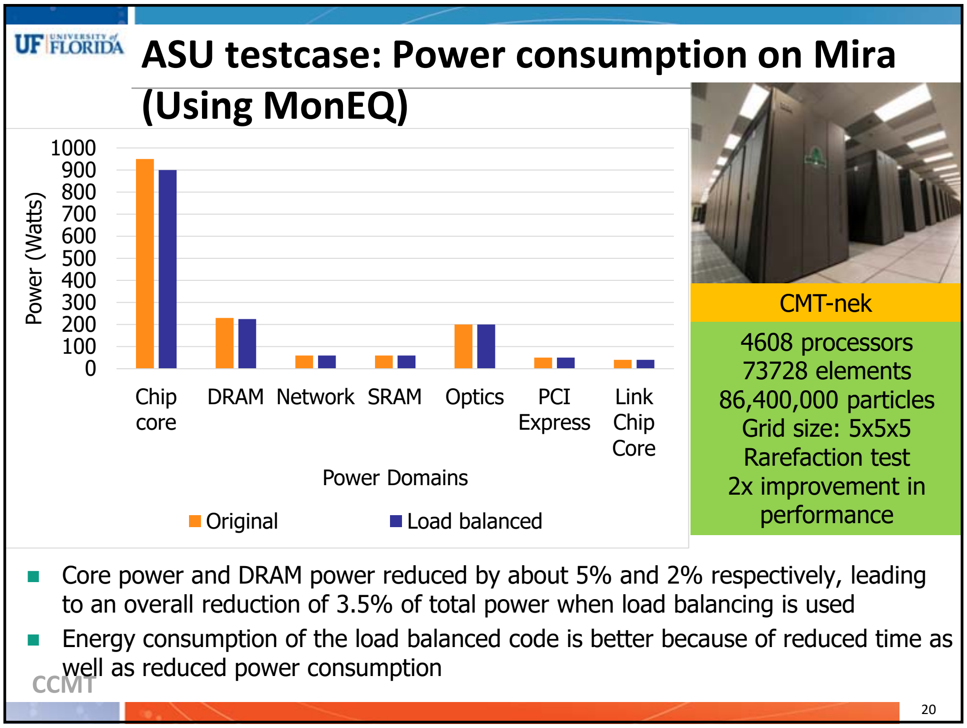

ASU testcase: Power consumption on Mira (Using MonEQ)

4608 processors
73728 elements
86,400,000 particles
Grid size: 5x5x5
Rarefaction test
2x improvement in performance

- Core power and DRAM power reduced by about 5% and 2% respectively, leading to an overall reduction of 3.5% of total power when load balancing is used
- Energy consumption of the load balanced code is better because of reduced time as well as reduced power consumption

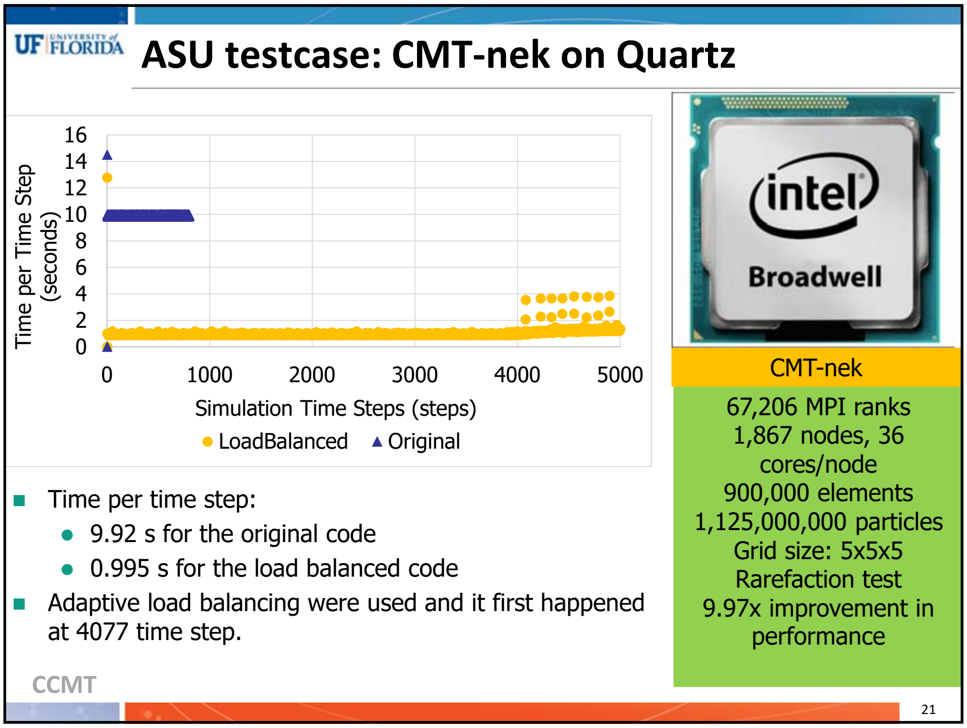

ASU testcase: CMT-nek on Quartz

- Time per time step:
  - 9.92 s for the original code
  - 0.995 s for the load balanced code
- Adaptive load balancing were used and it first happened at 4077 time step.

CMT-nek

67,206 MPI ranks
1,867 nodes, 36 cores/node
900,000 elements
1,125,000,000 particles
Grid size: 5x5x5
Rarefaction test
9.97x improvement in performance

CCMT

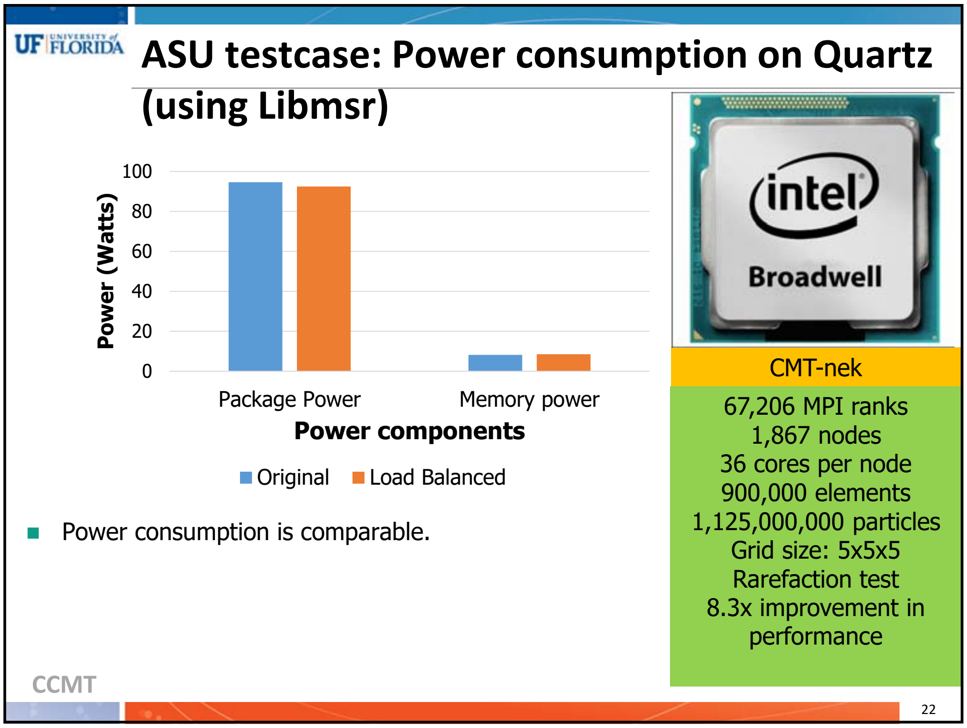

ASU testcase: Power consumption on Quartz (using Libmsr)

- Power consumption is comparable.

CMT-nek

67,206 MPI ranks
1,867 nodes
36 cores per node
900,000 elements
1,125,000,000 particles
Grid size: 5x5x5
Rarefaction test
8.3x improvement in performance

CCMT

Multilevel Memories

- What is a multi-level memory?
  - Main memory is composed of two or more types of memories

- Long term implications
  - Increase capacity
  - Decrease latency
  - Increase bandwidth
  - Decrease cost
  - Decrease power consumption
  - Challenges: Software and hardware challenges

- Why are they important for exascale
  - High bandwidth, high capacity, lower power consumption are attractive features

Multi-Level Memories

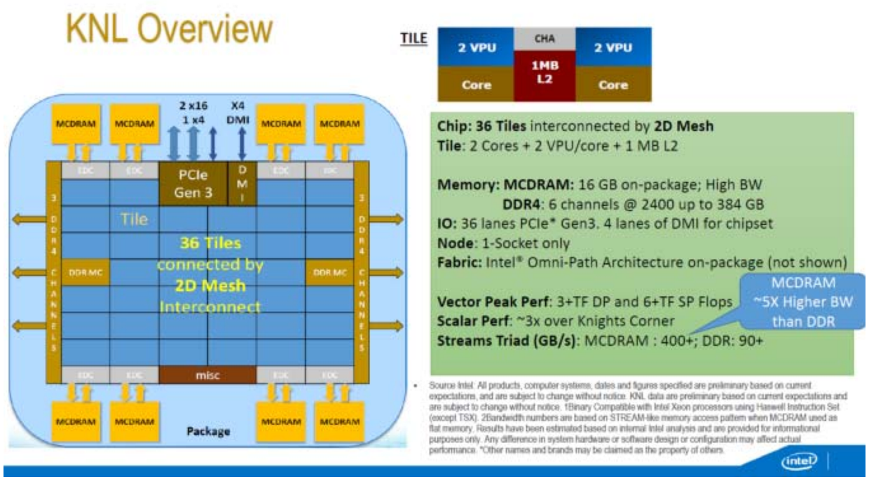

KNL Overview

Chip: 36 Tiles interconnected by **2D Mesh**
Tile: 2 Cores + 2 VPU/core + 1 MB L2

Memory: MCDRAM: 16 GB on-package; High BW
DDR4: 6 channels @ 2400 up to 384 GB
IO: 36 lanes PCIe* Gen3. 4 lanes of DMI for chipset
Node: 1-Socket only
Fabric: Intel® Omni-Path Architecture on-package (not shown)

Vector Peak Perf: 3+TF DP and 6+TF SP Flops
Scalar Perf: ~3x over Knights Corner
Streams Triad (GB/s): MCDRAM : 400+; DDR: 90+

MCDRAM ~5X Higher BW than DDR

- Source Intel: All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. KNL data are preliminary based on current expectations and are subject to change without notice. 1Binary Compatible with Intel Xeon processors using Haswell Instruction Set (except TSX). 2Bandwidth numbers are based on STREAM-like memory access pattern when MCDRAM used as flat memory. Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. *Other names and brands may be claimed as the property of others.

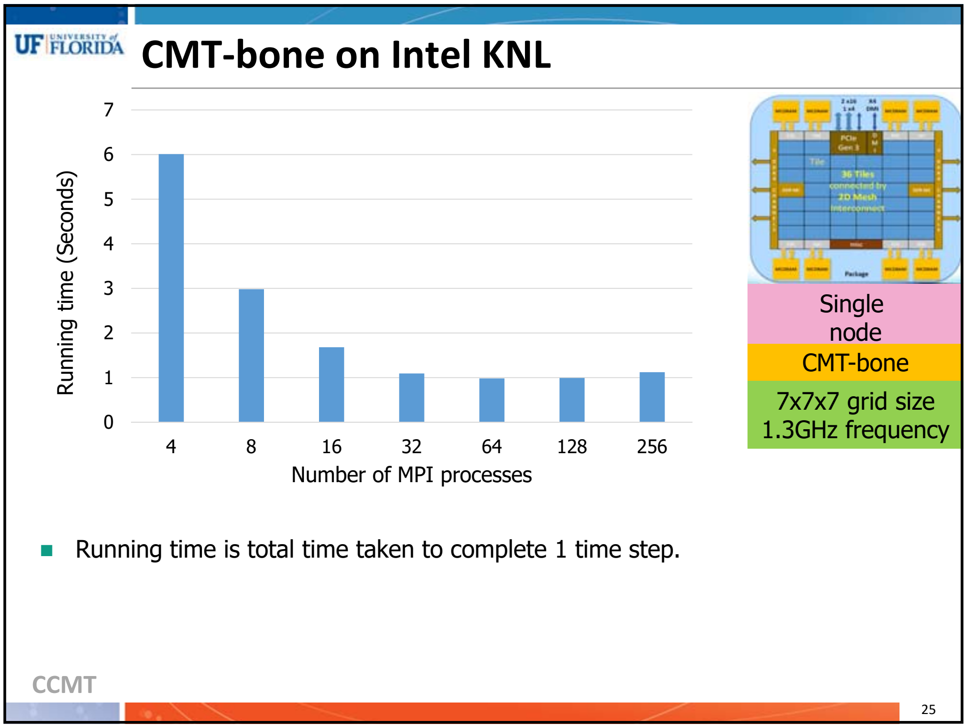

CMT-bone on Intel KNL

- Running time is total time taken to complete 1 time step.

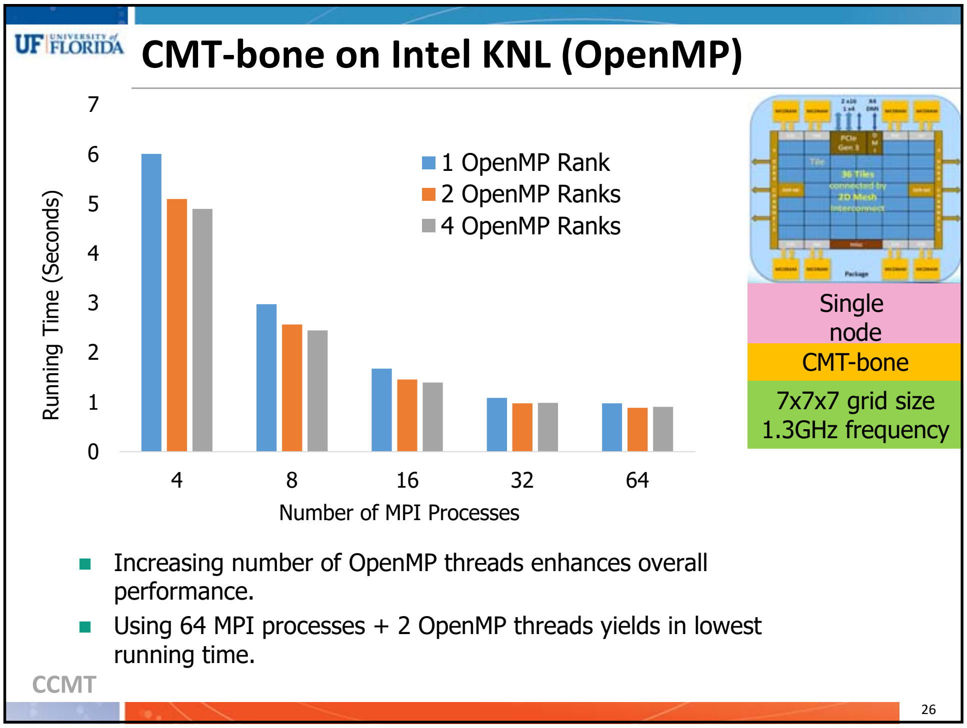

CMT-bone on Intel KNL (OpenMP)

- Increasing number of OpenMP threads enhances overall performance.
- Using 64 MPI processes + 2 OpenMP threads yields in lowest running time.

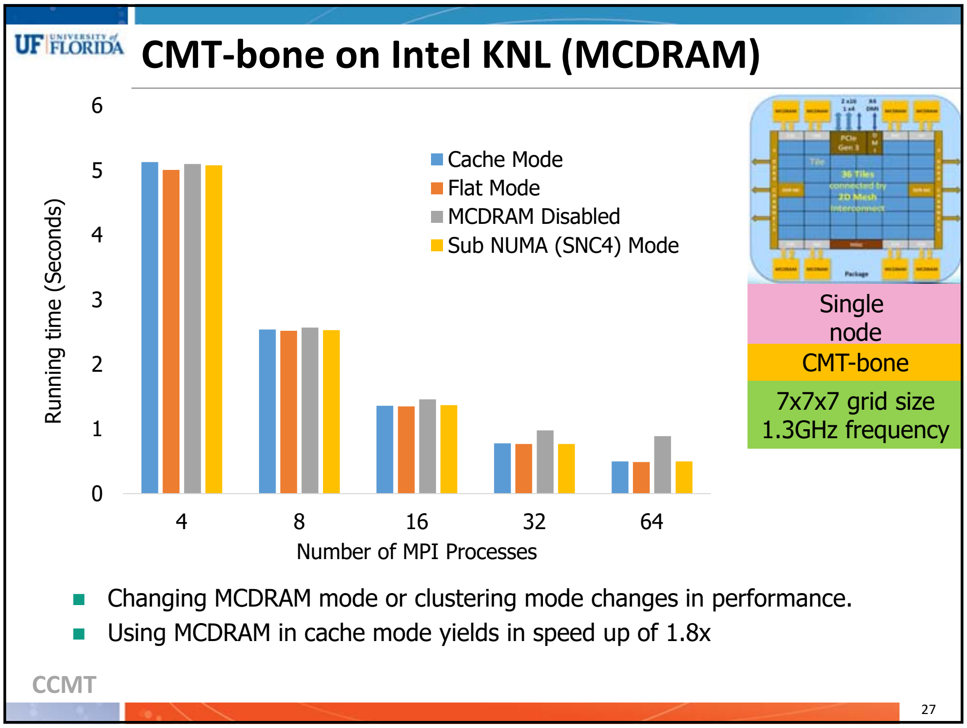

CMT-bone on Intel KNL (MCDRAM)

Single node

CMT-bone

7x7x7 grid size
1.3GHz frequency

- Changing MCDRAM mode or clustering mode changes in performance.
- Using MCDRAM in cache mode yields in speed up of 1.8x

CCMT

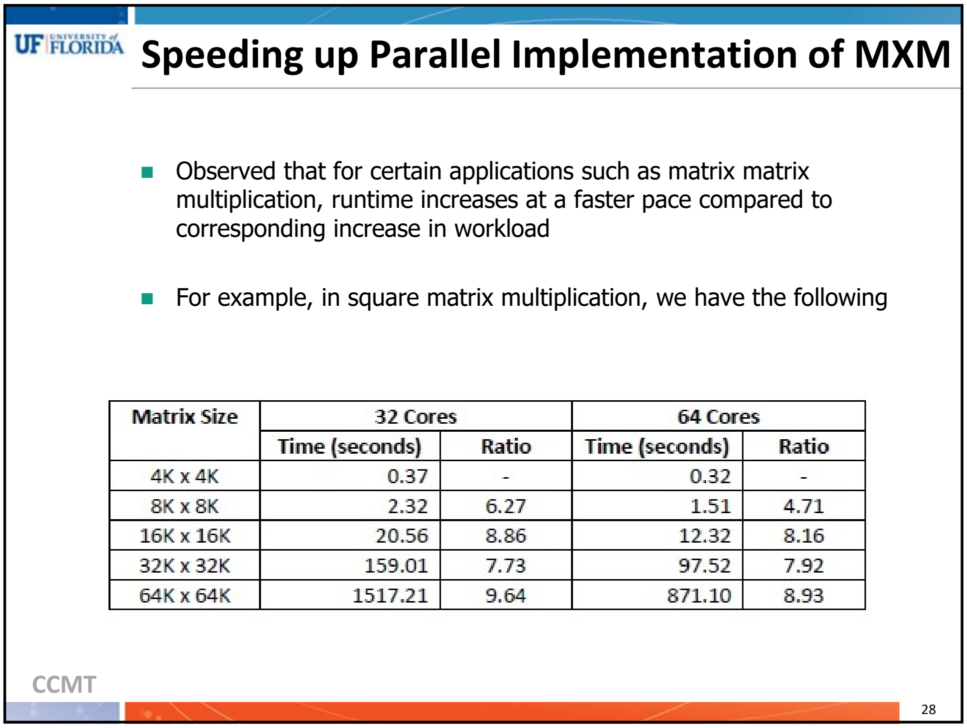

Speeding up Parallel Implementation of MXM

- Observed that for certain applications such as matrix matrix multiplication, runtime increases at a faster pace compared to corresponding increase in workload
- For example, in square matrix multiplication, we have the following

Matrix Size	32 Cores		64 Cores	
	Time (seconds)	Ratio	Time (seconds)	Ratio
4K x 4K	0.37	-	0.32	-
8K x 8K	2.32	6.27	1.51	4.71
16K x 16K	20.56	8.86	12.32	8.16
32K x 32K	159.01	7.73	97.52	7.92
64K x 64K	1517.21	9.64	871.10	8.93

CCMT

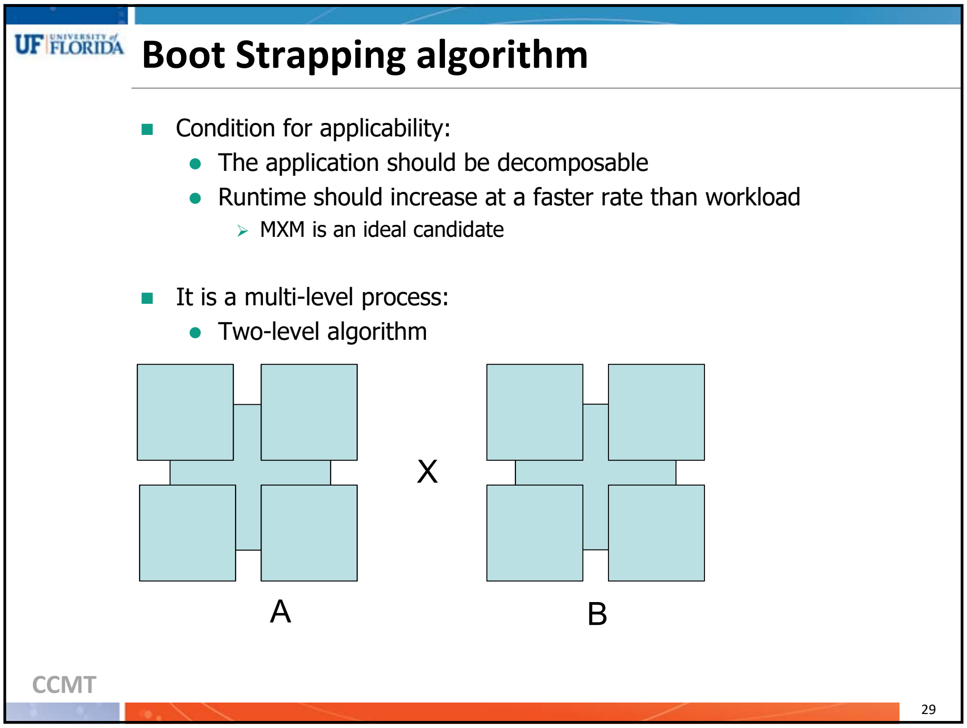

Boot Strapping algorithm

- Condition for applicability:
  - The application should be decomposable
  - Runtime should increase at a faster rate than workload
    - MXM is an ideal candidate
- It is a multi-level process:
  - Two-level algorithm

CCMT

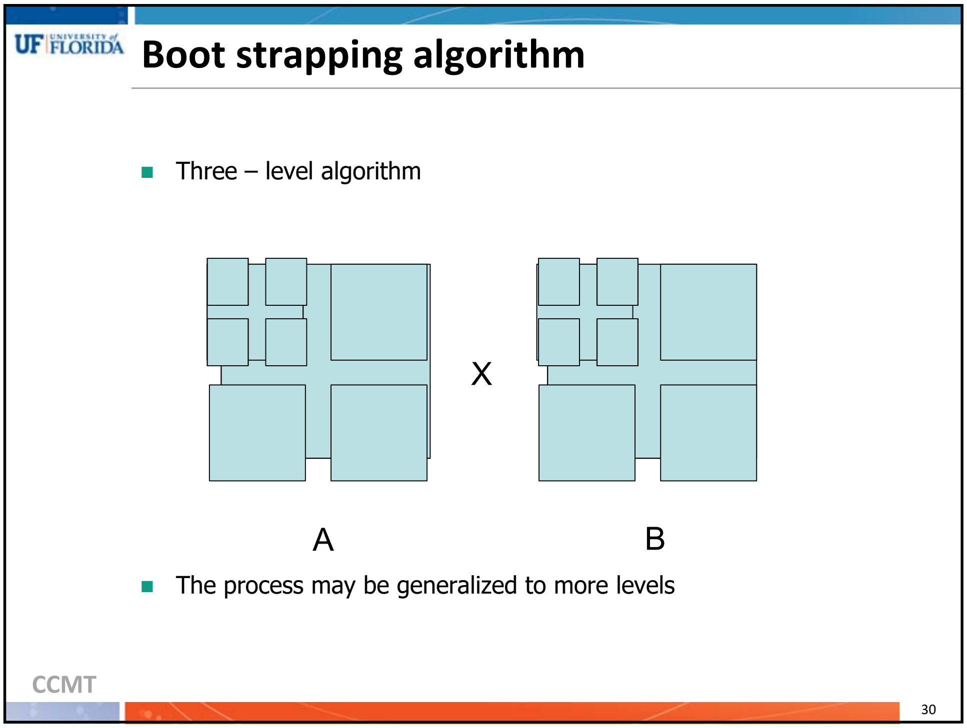

Boot strapping algorithm

- Three – level algorithm

- The process may be generalized to more levels

CCMT

Optimization to Reduce Cache Misses

- Matrix multiplication at the leaf level is done in an order that reduces cache misses

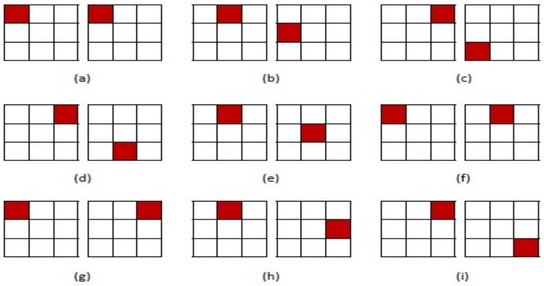

Theoretical Performance Analysis

- Using a simple performance model we can predict the expected speedup
- If T2 is the time taken by the two level algorithm, and T1 is the time taken by the original algorithm, then we can show that:

$$\frac{T_2}{T_1} = \frac{1+rbk\delta}{rbk\delta}$$

- Where, r = ratio of time taken to solve the original problem using all processors versus the time taken to solve the smallest decomposed matrix all the processors
- b = num allel iplications

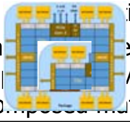

- k = ratio he ta the original problem using 1/b of the to er of versus the the time taken to solve the deco trix using the same number of processo

$r =$

$k =$

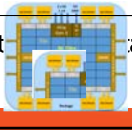

- delta = bination t aken to solve the original problem.

Experimental Setup

- Target application: Intel MKL cblas-dgemm on KNL platform
- Combinations of memory and clustering modes tested
- OpenMP parallel processing
- 64 cores
- Up to 16Kx16K : average over 10 runs
- 32Kx32K and 64Kx64K : average over 5 runs

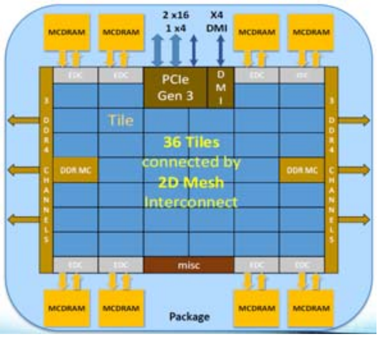

KNL Modes

- Memory modes:
  - Flat mode
  - Cache mode
  - Hybrid mode
- Clustering modes:
  - All-to-all
  - Quadrant
  - Hemisphere
  - SNC2
  - SNC4
- Clustering modes defines the affinity properties of tiles, tag directories and memory controllers

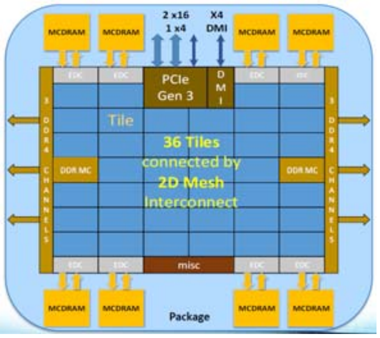

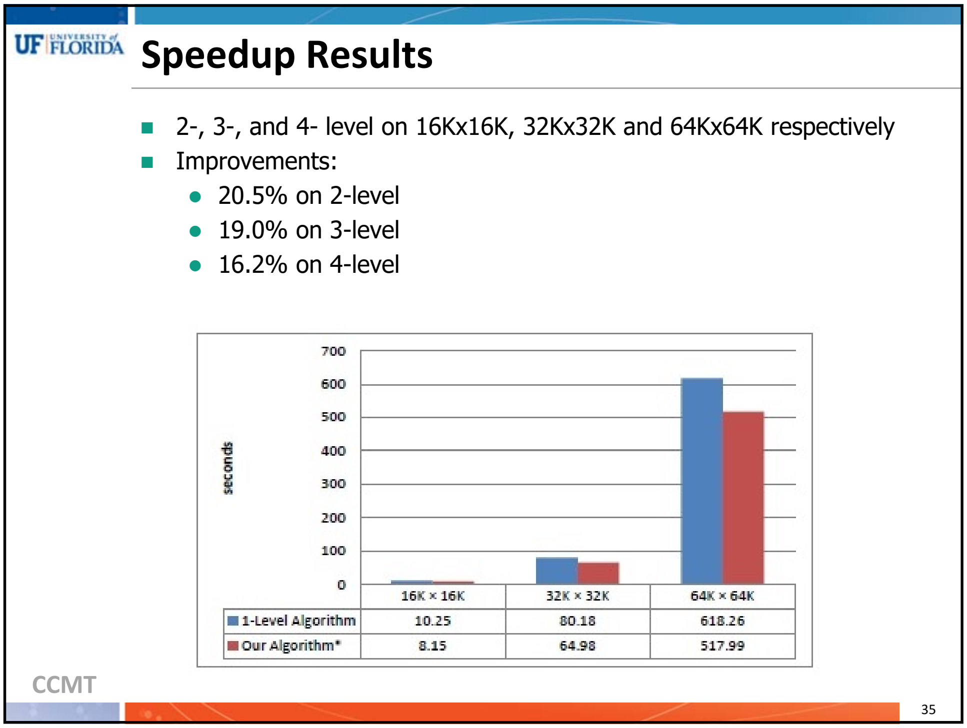

UF | UNIVERSITY of FLORIDA

Speedup Results

- 2-, 3-, and 4- level on 16Kx16K, 32Kx32K and 64Kx64K respectively
- Improvements:
  - 20.5% on 2-level
  - 19.0% on 3-level
  - 16.2% on 4-level

	16K × 16K	32K × 32K	64K × 64K
🟦 1-Level Algorithm	10.25	80.18	618.26
🟥 Our Algorithm*	8.15	64.98	517.99

CCMT

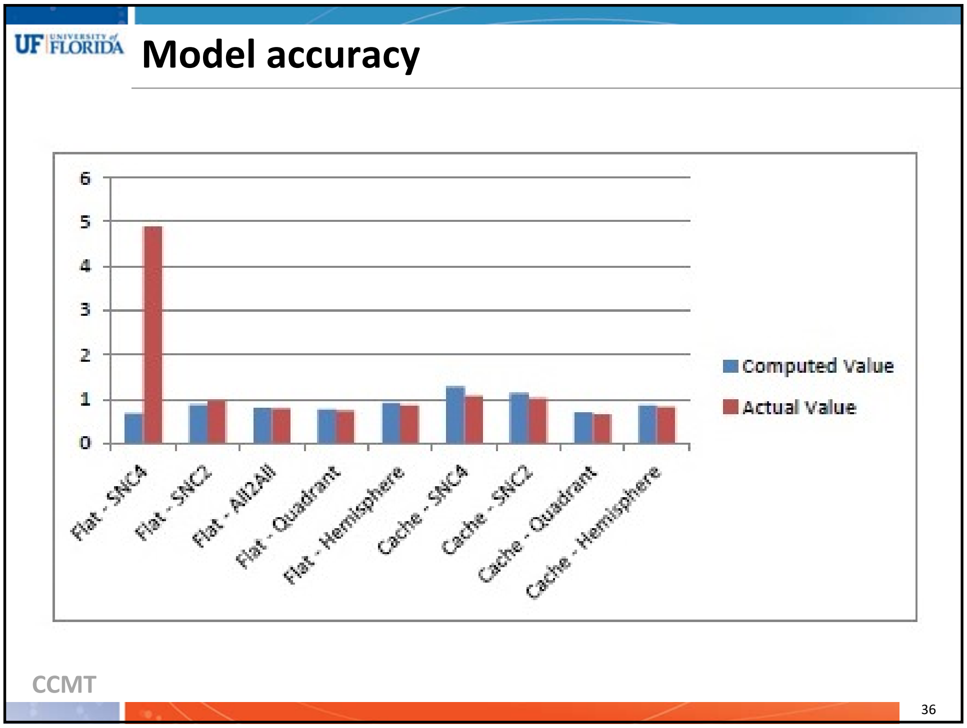

UF | UNIVERSITY of FLORIDA

Model accuracy

CCMT

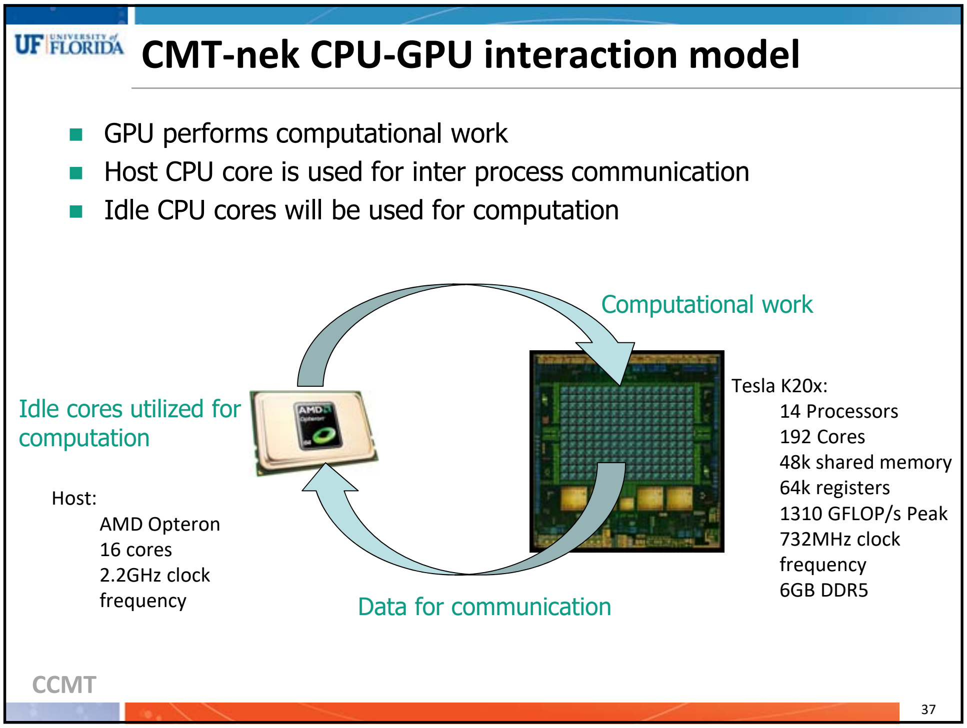

CMT-nek CPU-GPU interaction model

- GPU performs computational work
- Host CPU core is used for inter process communication
- Idle CPU cores will be used for computation

Computational work

Idle cores utilized for computation

Host:
AMD Opteron
16 cores
2.2GHz clock
frequency

Data for communication

Tesla K20x:
14 Processors
192 Cores
48k shared memory
64k registers
1310 GFLOP/s Peak
732MHz clock
frequency
6GB DDR5

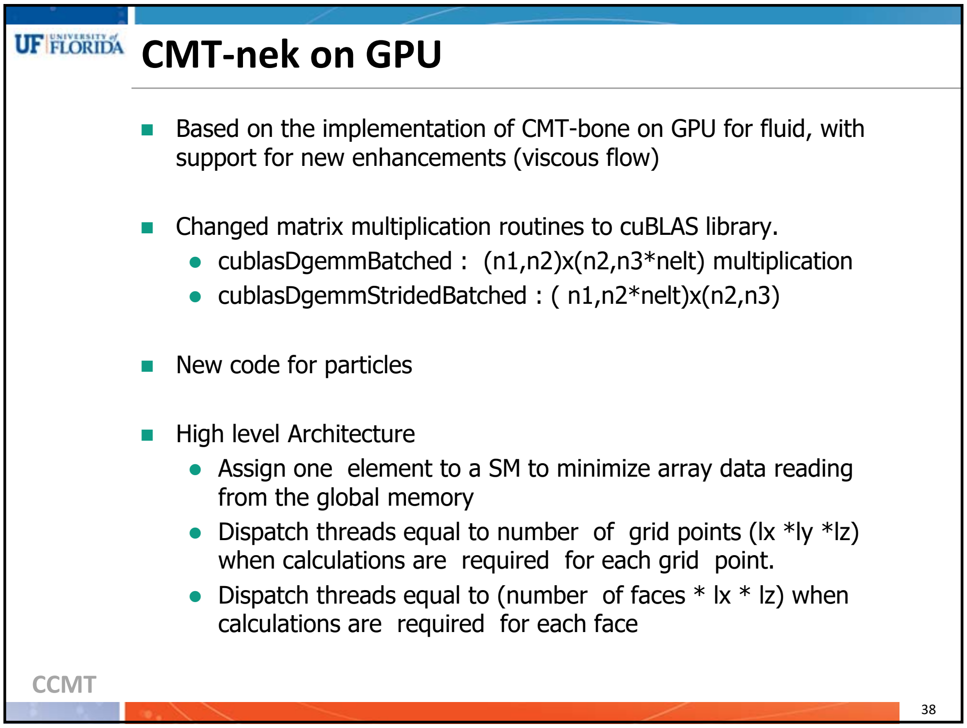

CMT-nek on GPU

- Based on the implementation of CMT-bone on GPU for fluid, with support for new enhancements (viscous flow)
- Changed matrix multiplication routines to cuBLAS library.
  - cublasDgemmBatched : (n1,n2)x(n2,n3*nelt) multiplication
  - cublasDgemmStridedBatched : (n1,n2*nelt)x(n2,n3)
- New code for particles
- High level Architecture
  - Assign one element to a SM to minimize array data reading from the global memory
  - Dispatch threads equal to number of grid points (lx *ly *lz) when calculations are required for each grid point.
  - Dispatch threads equal to (number of faces * lx * lz) when calculations are required for each face

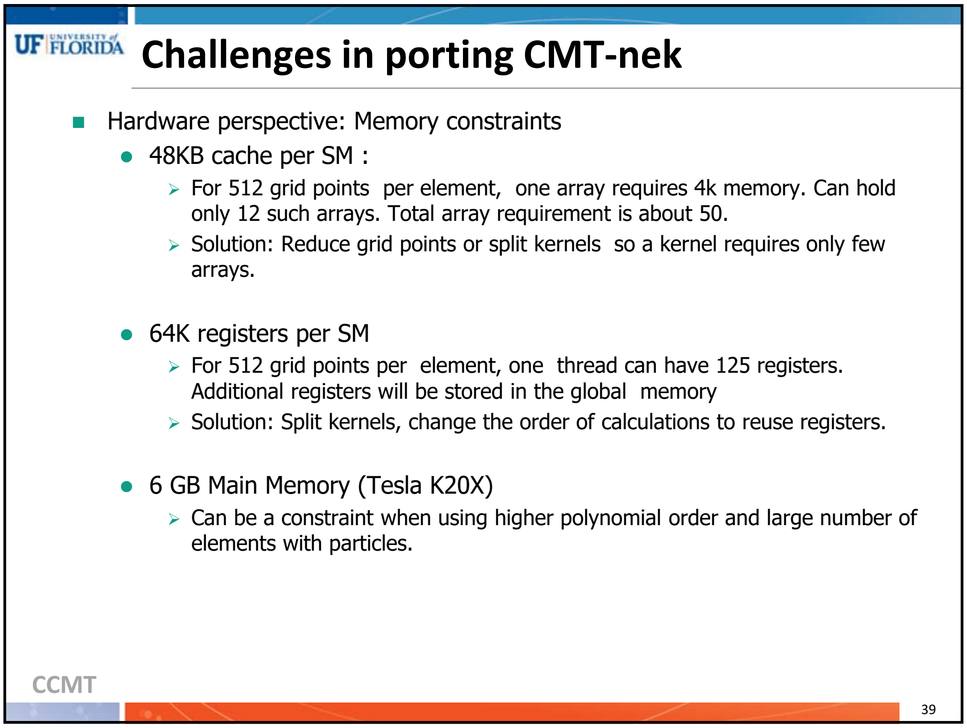

Challenges in porting CMT-nek

- Hardware perspective: Memory constraints
  - 48KB cache per SM :
    - For 512 grid points per element, one array requires 4k memory. Can hold only 12 such arrays. Total array requirement is about 50.
    - Solution: Reduce grid points or split kernels so a kernel requires only few arrays.
  - 64K registers per SM
    - For 512 grid points per element, one thread can have 125 registers. Additional registers will be stored in the global memory
    - Solution: Split kernels, change the order of calculations to reuse registers.
  - 6 GB Main Memory (Tesla K20X)
    - Can be a constraint when using higher polynomial order and large number of elements with particles.

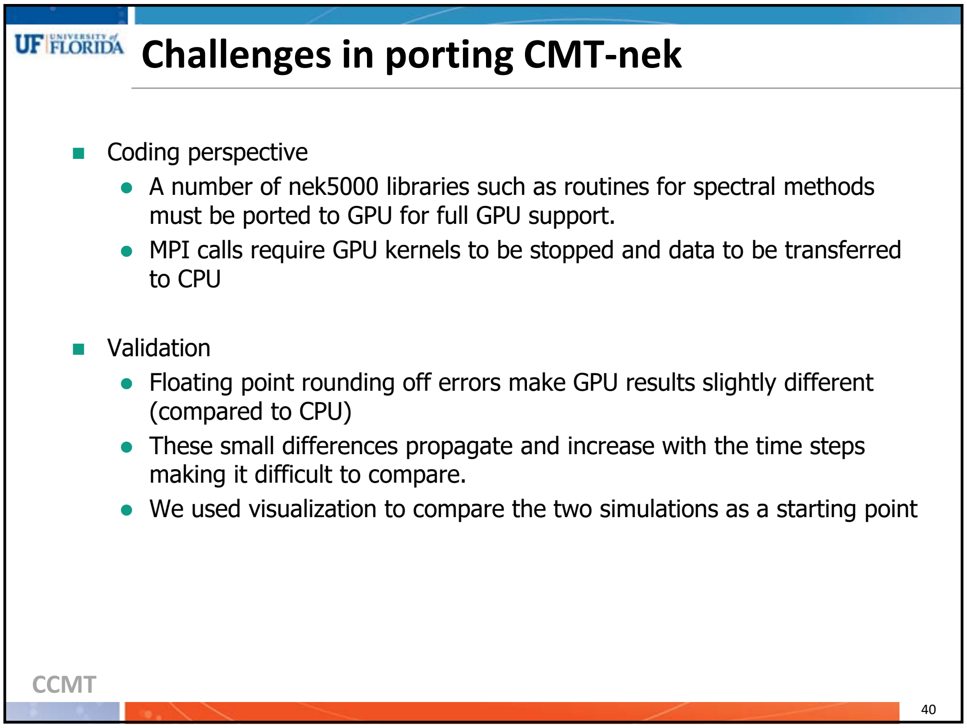

Challenges in porting CMT-nek

- Coding perspective
  - A number of nek5000 libraries such as routines for spectral methods must be ported to GPU for full GPU support.
  - MPI calls require GPU kernels to be stopped and data to be transferred to CPU
- Validation
  - Floating point rounding off errors make GPU results slightly different (compared to CPU)
  - These small differences propagate and increase with the time steps making it difficult to compare.
  - We used visualization to compare the two simulations as a starting point

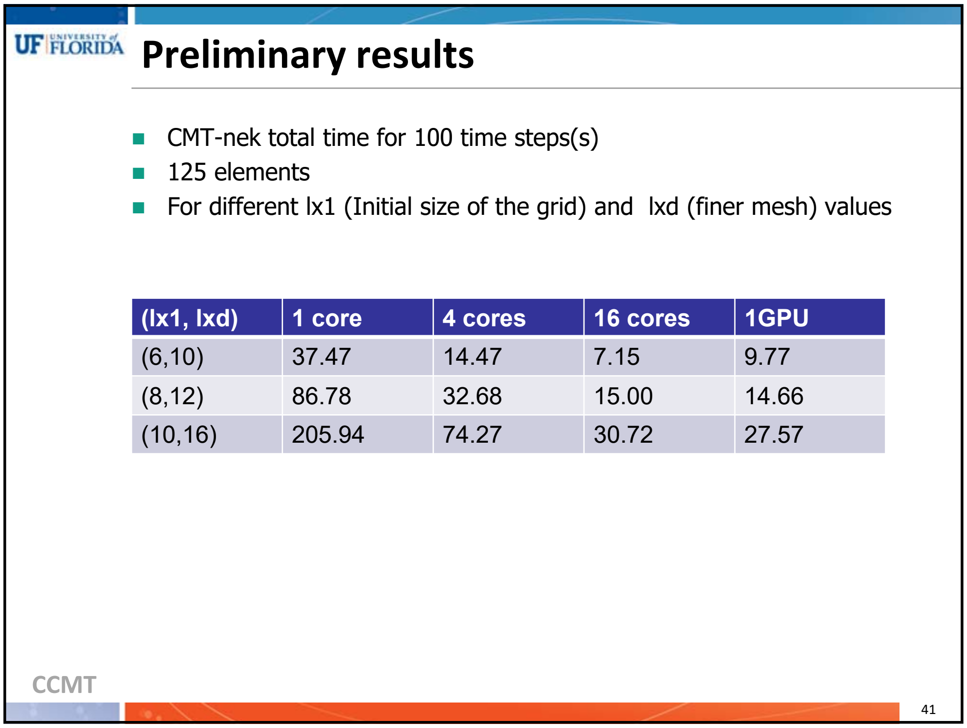

Preliminary results

- CMT-nek total time for 100 time steps(s)
- 125 elements
- For different lx1 (Initial size of the grid) and lxd (finer mesh) values

(lx1, lxd)	1 core	4 cores	16 cores	1GPU
(6,10)	37.47	14.47	7.15	9.77
(8,12)	86.78	32.68	15.00	14.66
(10,16)	205.94	74.27	30.72	27.57

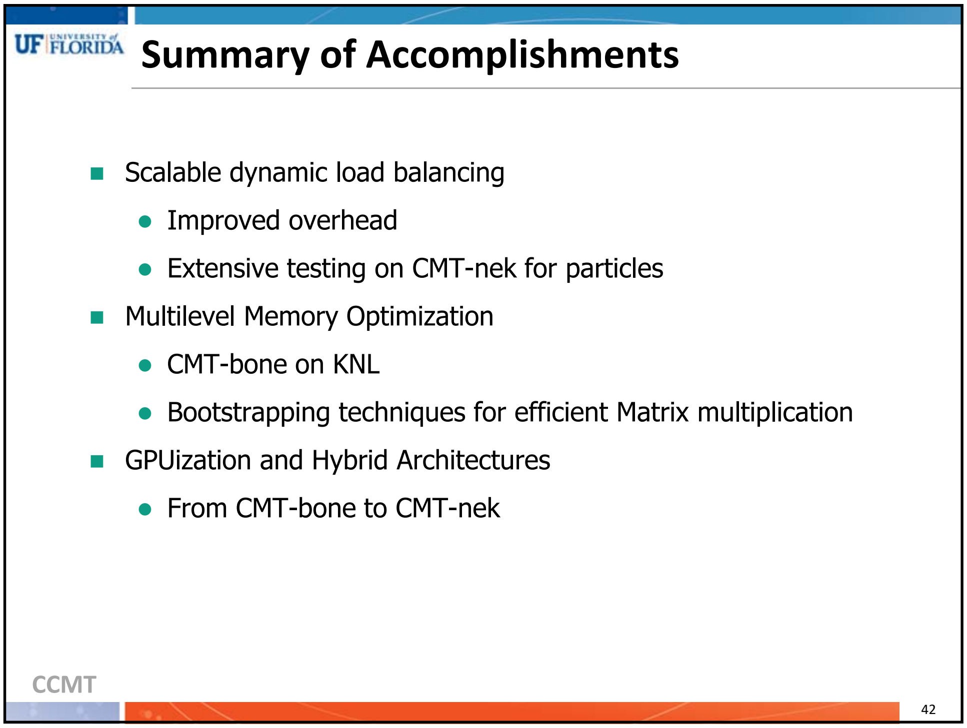

Summary of Accomplishments

- Scalable dynamic load balancing
  - Improved overhead
  - Extensive testing on CMT-nek for particles
- Multilevel Memory Optimization
  - CMT-bone on KNL
  - Bootstrapping techniques for efficient Matrix multiplication
- GPUization and Hybrid Architectures
  - From CMT-bone to CMT-nek

UF UNIVERSITY of FLORIDA

Publications

1. Tania Banerjee and Sanjay Ranka, Genetic Algorithm based Autotuning Approach for Performance and Energy Optimization, Proceedings of International Green and Sustainable Computing, 2015
2. Tania Banerjee, Mohammed Gadou, Sanjay Ranka, A Genetic Algorithm based Approach for Multi-objective Hardware/Software Co-optimization, Journal of Sustainable Computing, Vol 10, June 2016, 36 – 47.
3. Tania Banerjee, Jacob Rabb and Sanjay Ranka, Multi-objective Optimization of Spectral Solvers on Hybrid Multicore Platforms, Journal of Sustainable Computing, Vol 12, Dec 2016, 10 - 20.
4. Tania Banerjee, Jason Hackl, Mrugesh Sringarpure, Tanzima Islam, S. Balachandar, Thomas Jackson, Sanjay Ranka, CMT-bone -- A Proxy Application for Compressible Multiphase Turbulent Flows, Proceedings of HiPC, 2016
5. Mohamed Gadou, Tania Banerjee, Sanjay Ranka, Multi-objective Optimization of CMT-bone on Hybrid Processors, Proceedings of IGSC 2016
6. Tania Banerjee, Jason Hackl, Mrugesh Sringarpure, Tanzima Islam, S. Balachandar, Thomas Jackson, Sanjay Ranka, A New Proxy Application for Compressible Multiphase Turbulent Flows. Sustainable Computing: Informatics and Systems, 2017
7. Mohamed Gadou, Tania Banerjee, Meena Arunachalam, Sanjay Ranka, Multiobjective evaluation and optimization of CMT-bone on multiple CPU/GPU systems, Journal of Sustainable Computing: Informatics and Systems, 2017
8. Keke Zhai, Tania Banerjee, David Zwick, Jason Hackl and Sanjay Ranka, Dynamic Load Balancing for Compressible Multiphase Turbulence, accepted International Conference on Supercomputing, 2018
9. Mohamed Gadou, Sankeerth Mogili, Tania Banerjee, Sanjay Ranka, Multi-objective Optimization on Hybrid Systems using DVFS, in preparation
10. Mohamed Gadou, Meena Arunachalam, Tania Banerjee, Sanjay Ranka, Multiobjective evaluation and optimization of CMT-bone on Intel Knights Landing, accepted in IGSC 2018.

43

CCMT

Do you have any questions?

UF UNIVERSITY of FLORIDA

NNSA

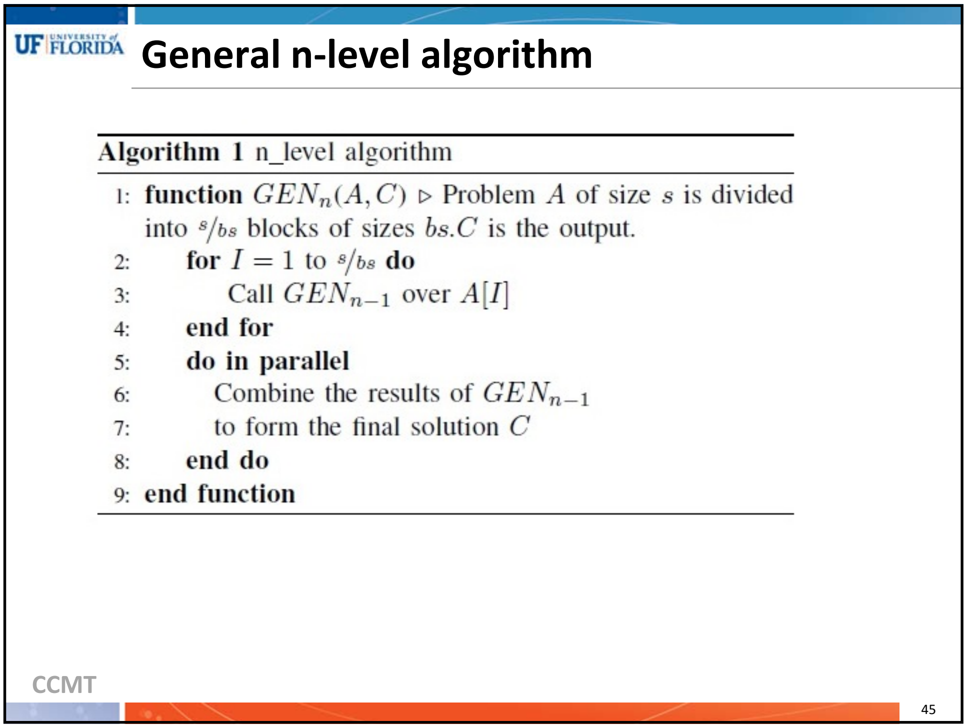

General n-level algorithm

Algorithm 1 n_level algorithm

```
function $GEN_n(A, C)$ ▷ Problem $A$ of size $s$ is divided into $s/bs$ blocks of sizes $bs$.$C$ is the output.
    for $I = 1$ to $s/bs$ do
        Call $GEN_{n-1}$ over $A[I]$
    end for
    do in parallel
        Combine the results of $GEN_{n-1}$
        to form the final solution $C$
    end do
end function
```

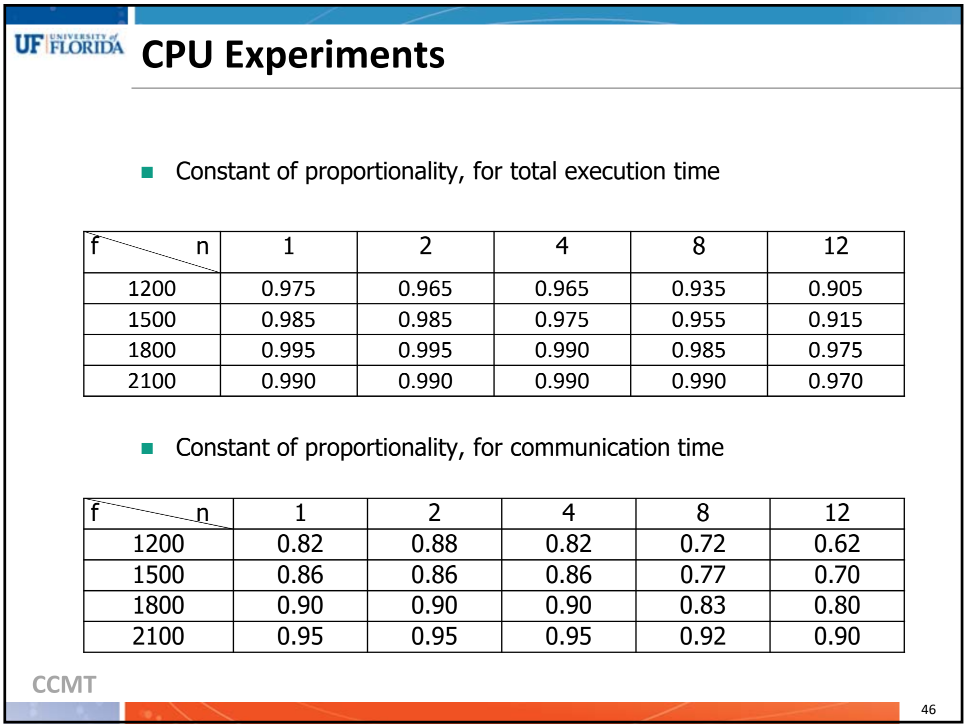

CPU Experiments

- Constant of proportionality, for total execution time

f n	1	2	4	8	12
1200	0.975	0.965	0.965	0.935	0.905
1500	0.985	0.985	0.975	0.955	0.915
1800	0.995	0.995	0.990	0.985	0.975
2100	0.990	0.990	0.990	0.990	0.970

- Constant of proportionality, for communication time

f n	1	2	4	8	12
1200	0.82	0.88	0.82	0.72	0.62
1500	0.86	0.86	0.86	0.77	0.70
1800	0.90	0.90	0.90	0.83	0.80
2100	0.95	0.95	0.95	0.92	0.90

Hardware-Software Autotuning

```
Algorithm: dudr-4loop
do k = 1, $N_z$
  do j = 1, $N_y$
    do i = 1, $N_x$
      do l = 1, $N_x$
        dudr(l, j, k) = dudr(l, j, k) +
                        a(i, l) * u(l, j, k, ie)
      enddo
    enddo
  enddo
enddo
```

```
Algorithm: dudr-4loop-fused
do k = 1, $N_z$* $N_y$
  do i = 1, $N_x$
    do l = 1, $N_x$
      dudr(l, k) = dudr(l, k) +
                   a(i, l) * u(l, k, ie)
    enddo
  enddo
enddo
```

```
Algorithm: dudr-4loop-permuted-and unrolled
do k = 1, $N_z$
 do i = 1, $N_x$
  do j = 1, $N_y$
   do l = 1, $N_x$,2
    dudr(l, j, k) = dudr(l, j, k) +
                    a(i, l) * u(l, j, k, ie)
    dudr(l+1, j, k) = dudr(l+1, j, k) +
                    a(i, l+1) * u(l+1, j, k, ie)
   enddo
  enddo
 enddo
enddo
```

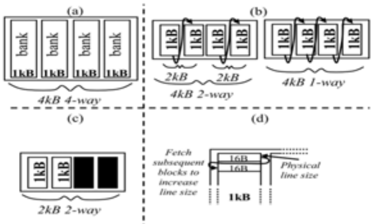

Total number of variants = 20

Number of 4-loop implementations for dudr

$N_x=N_y=N_z=10$

= 4! * 4 ^ 4

= 24 * 256 = 6144 variants

Total number of variants = 98,240 (N=10)

Total number of variants = 217,728 (N=20)

Exhaustive search may not be feasible

Genetic Algorithm

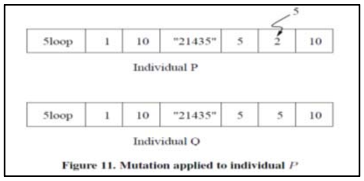

Figure 11. Mutation applied to individual *P*

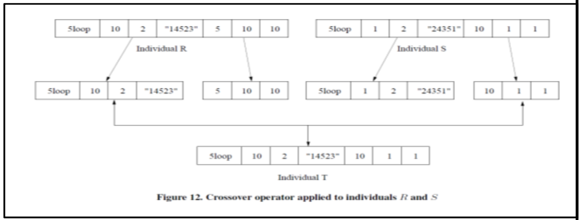

Figure 12. Crossover operator applied to individuals *R* and *S*

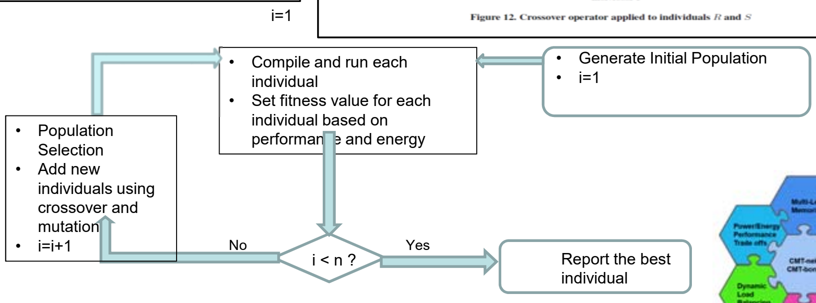

i=1

- Compile and run each individual
- Set fitness value for each individual based on performance and energy

- Generate Initial Population
- i=1

- Population Selection
- Add new individuals using crossover and mutation
- i=i+1

No

i < n ?

Yes

Report the best individual

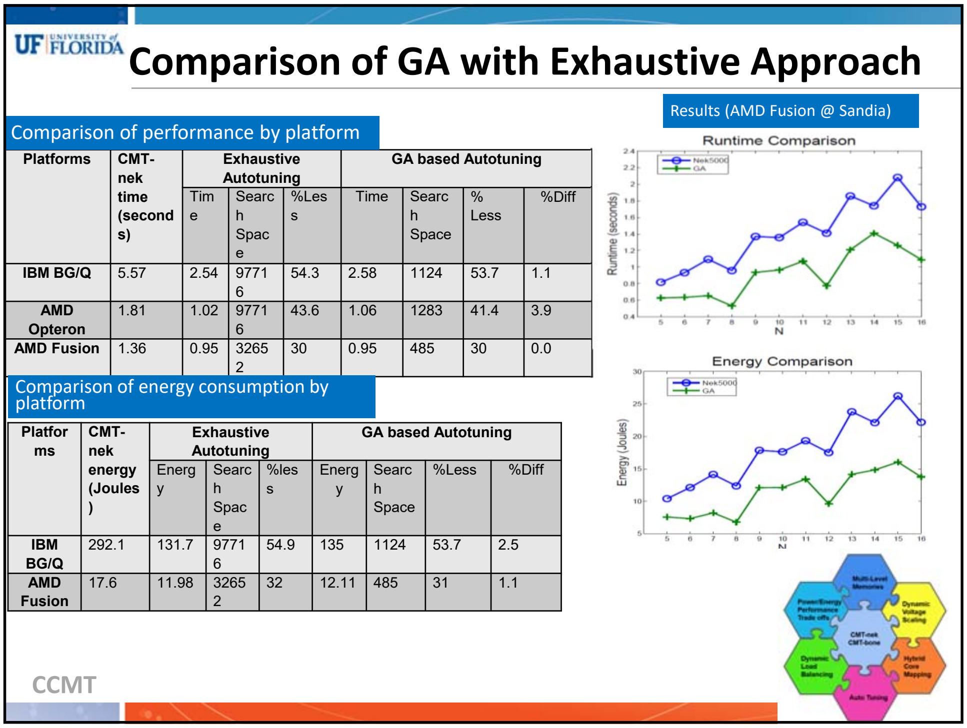

Comparison of GA with Exhaustive Approach

Comparison of performance by platform

Platforms	CMT-nek time (seconds)	Exhaustive Autotuning			GA based Autotuning			
		Time	Search Space	%Less	Time	Search Space	% Less	%Diff
IBM BG/Q	5.57	2.54	97716	54.3	2.58	1124	53.7	1.1
AMD Opteron	1.81	1.02	97716	43.6	1.06	1283	41.4	3.9
AMD Fusion	1.36	0.95	32652	30	0.95	485	30	0.0

Comparison of energy consumption by platform

Platforms	CMT-nek energy (Joules)	Exhaustive Autotuning			GA based Autotuning			
		Energy	Search Space	%less	Energy	Search Space	%Less	%Diff
IBM BG/Q	292.1	131.7	97716	54.9	135	1124	53.7	2.5
AMD Fusion	17.6	11.98	32652	32	12.11	485	31	1.1

Results (AMD Fusion @ Sandia)

CCMT

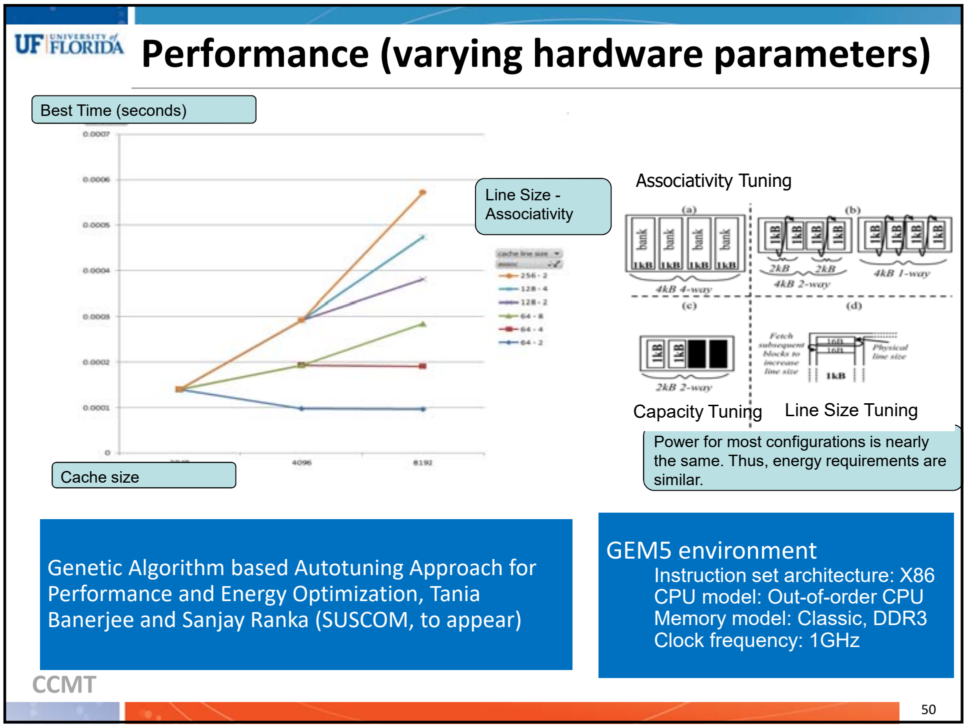

Performance (varying hardware parameters)

Best Time (seconds)

Cache size

Line Size - Associativity

Associativity Tuning

Capacity Tuning

Line Size Tuning

Power for most configurations is nearly the same. Thus, energy requirements are similar.

Genetic Algorithm based Autotuning Approach for Performance and Energy Optimization, Tania Banerjee and Sanjay Ranka (SUSCOM, to appear)

GEM5 environment
- Instruction set architecture: X86
- CPU model: Out-of-order CPU
- Memory model: Classic, DDR3
- Clock frequency: 1GHz

CCMT

CCMT

Exascale Behavioral Emulation

Principal Investigators:

Herman Lam, Greg Stitt
Center for Compressible Multiphase Turbulence (CCMT)
NSF Center for High-Performance Reconfigurable Computing (CHREC)
ECE Department, University of Florida

Exascale Behavioral Emulation (BE) Team


Ryan Blanchard
FPGA Acceleration


Sai Chenna
BE Design Space Exp.


Trokon Johnson
BE Tools


Aravind Neelakantan
BE Methods


Carlo Pascoe
FPGA Acceleration


Maneesh Merugu
BE-SST


Steven Paek
Thermal Modeling

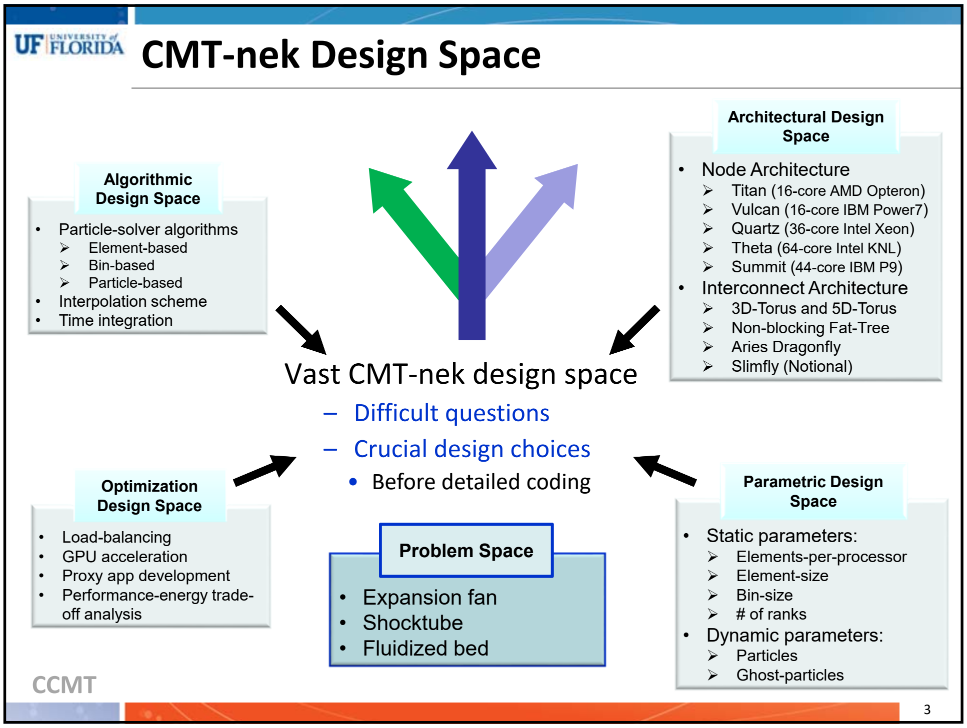

CMT-nek Design Space

Algorithmic Design Space
- Particle-solver algorithms
  - Element-based
  - Bin-based
  - Particle-based
- Interpolation scheme
- Time integration

Architectural Design Space
- Node Architecture
  - Titan (16-core AMD Opteron)
  - Vulcan (16-core IBM Power7)
  - Quartz (36-core Intel Xeon)
  - Theta (64-core Intel KNL)
  - Summit (44-core IBM P9)
- Interconnect Architecture
  - 3D-Torus and 5D-Torus
  - Non-blocking Fat-Tree
  - Aries Dragonfly
  - Slimfly (Notional)

Vast CMT-nek design space
- Difficult questions
- Crucial design choices
  - Before detailed coding

Optimization Design Space
- Load-balancing
- GPU acceleration
- Proxy app development
- Performance-energy trade-off analysis

Problem Space
- Expansion fan
- Shocktube
- Fluidized bed

Parametric Design Space
- Static parameters:
  - Elements-per-processor
  - Element-size
  - Bin-size
  - # of ranks
- Dynamic parameters:
  - Particles
  - Ghost-particles

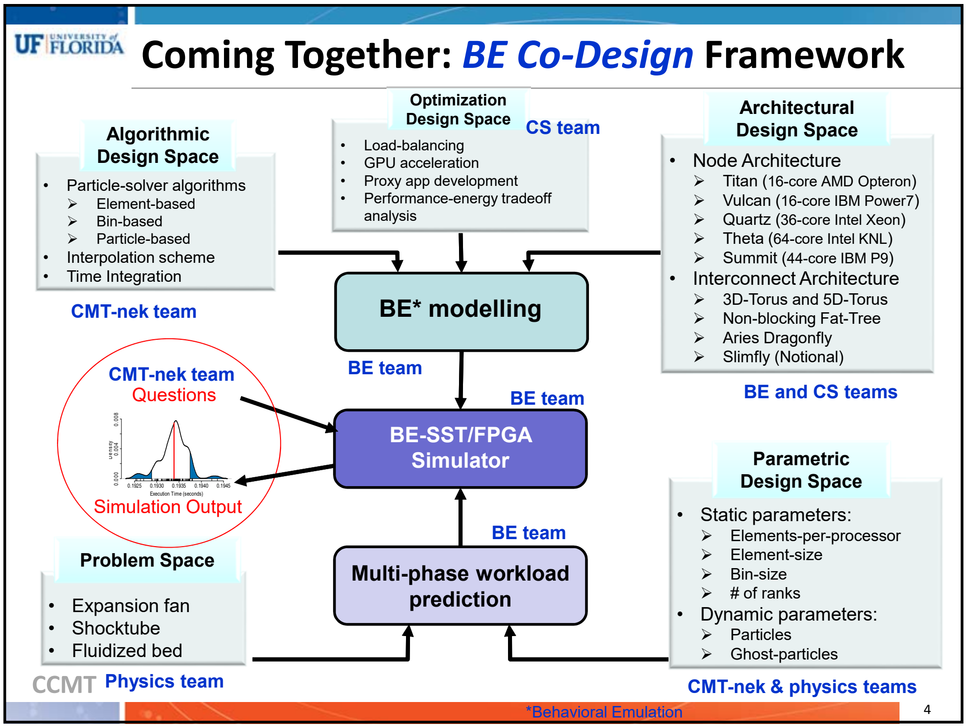

Coming Together: *BE Co-Design* Framework

Algorithmic Design Space
- Particle-solver algorithms
  - Element-based
  - Bin-based
  - Particle-based
- Interpolation scheme
- Time Integration

CMT-nek team

Optimization Design Space CS team
- Load-balancing
- GPU acceleration
- Proxy app development
- Performance-energy tradeoff analysis

Architectural Design Space
- Node Architecture
  - Titan (16-core AMD Opteron)
  - Vulcan (16-core IBM Power7)
  - Quartz (36-core Intel Xeon)
  - Theta (64-core Intel KNL)
  - Summit (44-core IBM P9)
- Interconnect Architecture
  - 3D-Torus and 5D-Torus
  - Non-blocking Fat-Tree
  - Aries Dragonfly
  - Slimfly (Notional)

BE and CS teams

BE* modelling

BE team

CMT-nek team
Questions
Simulation Output

BE team

BE-SST/FPGA Simulator

BE team

Multi-phase workload prediction

Problem Space
- Expansion fan
- Shocktube
- Fluidized bed

Physics team

Parametric Design Space
- Static parameters:
  - Elements-per-processor
  - Element-size
  - Bin-size
  - # of ranks
- Dynamic parameters:
  - Particles
  - Ghost-particles

CMT-nek & physics teams

*Behavioral Emulation

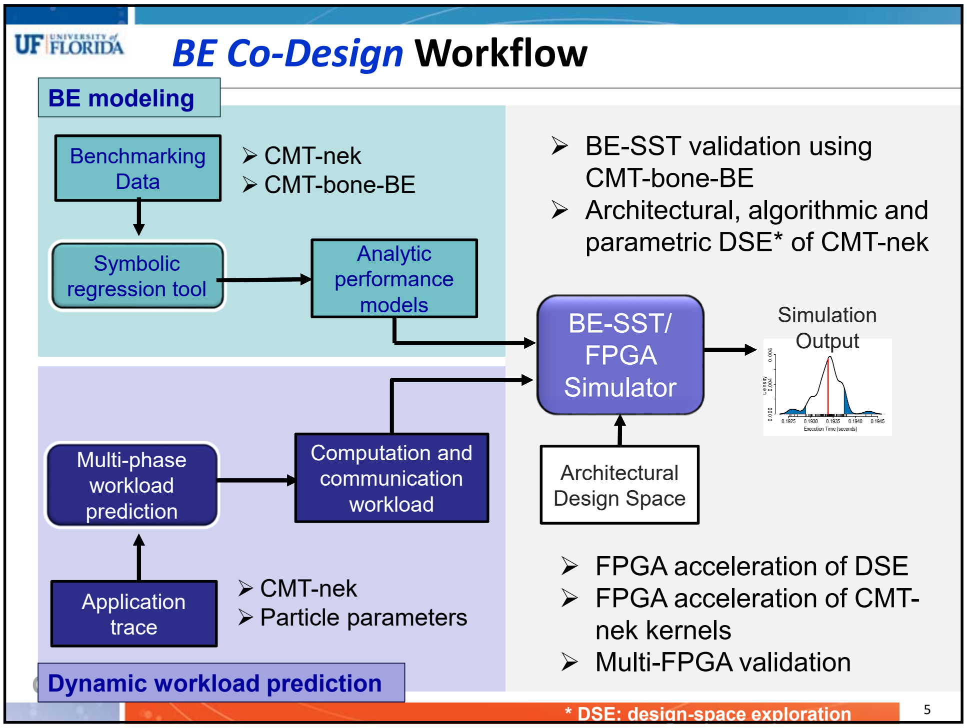

BE Co-Design Workflow

BE modeling

Benchmarking Data → Symbolic regression tool → Analytic performance models → BE-SST/ FPGA Simulator

- CMT-nek
- CMT-bone-BE

- BE-SST validation using CMT-bone-BE
- Architectural, algorithmic and parametric DSE* of CMT-nek

BE-SST/ FPGA Simulator → Simulation Output

Architectural Design Space → BE-SST/ FPGA Simulator

- FPGA acceleration of DSE
- FPGA acceleration of CMT-nek kernels
- Multi-FPGA validation

Application trace → Multi-phase workload prediction → Computation and communication workload → BE-SST/ FPGA Simulator

- CMT-nek
- Particle parameters

Dynamic workload prediction

* DSE: design-space exploration

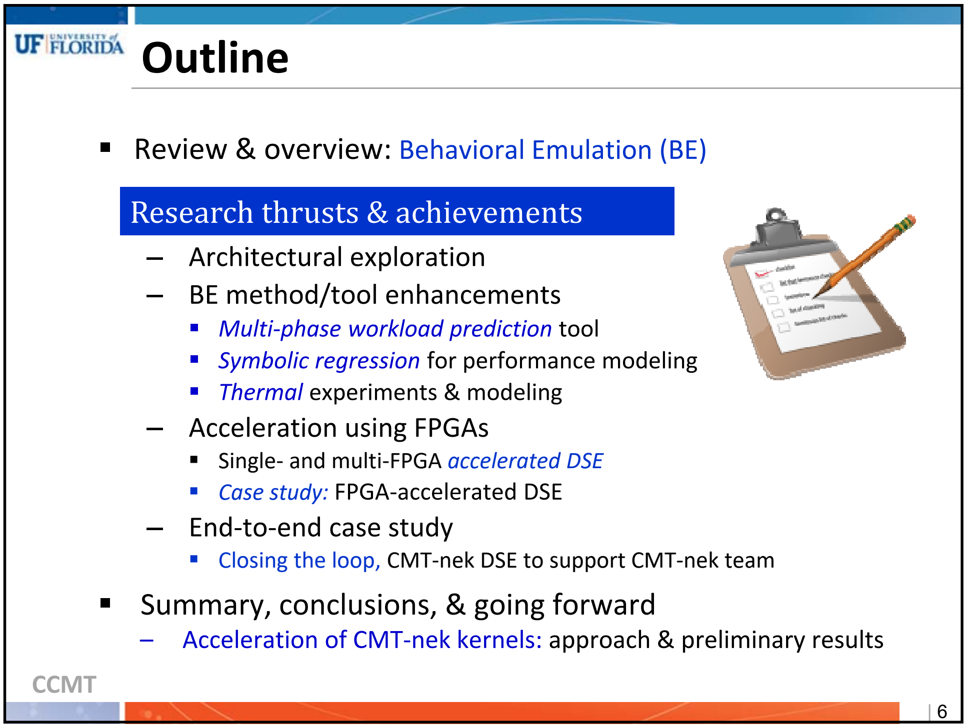

Outline

- Review & overview: Behavioral Emulation (BE)

Research thrusts & achievements

- Architectural exploration
- BE method/tool enhancements
  - *Multi-phase workload prediction* tool
  - *Symbolic regression* for performance modeling
  - *Thermal* experiments & modeling
- Acceleration using FPGAs
  - Single- and multi-FPGA *accelerated DSE*
  - *Case study:* FPGA-accelerated DSE
- End-to-end case study
  - Closing the loop, CMT-nek DSE to support CMT-nek team

- Summary, conclusions, & going forward
  - Acceleration of CMT-nek kernels: approach & preliminary results

CCMT

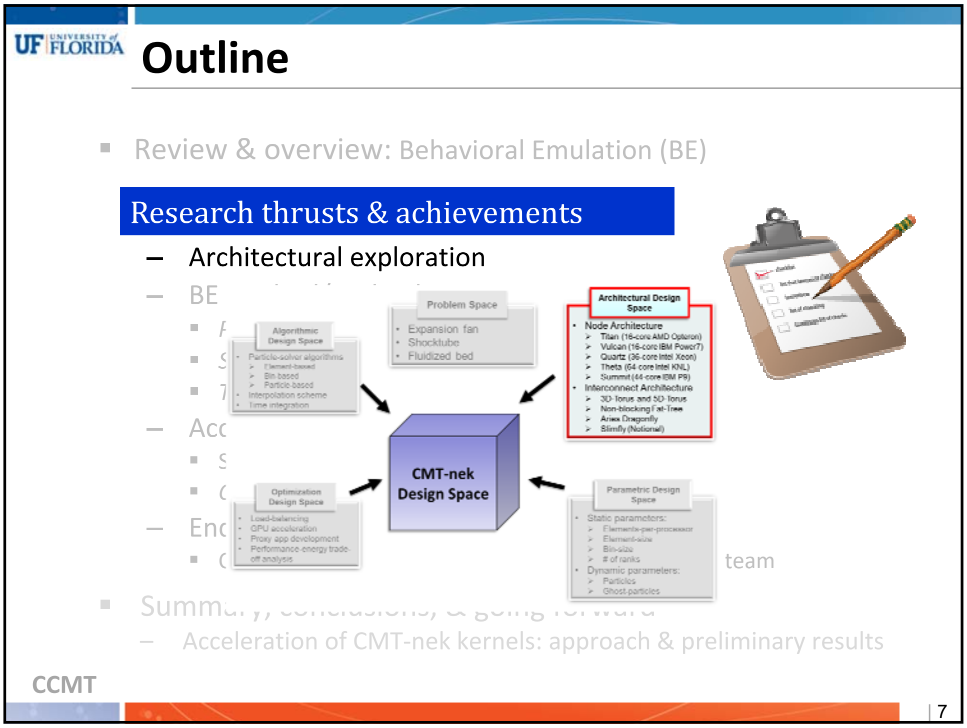

Outline

- Review & overview: Behavioral Emulation (BE)

Research thrusts & achievements

- Architectural exploration
- BE
  - F
  - S
  - T
- Acc
  - S
  - C
- Enc
  - C team
- Summ… , conclusions, & going forward
  - Acceleration of CMT-nek kernels: approach & preliminary results

CCMT

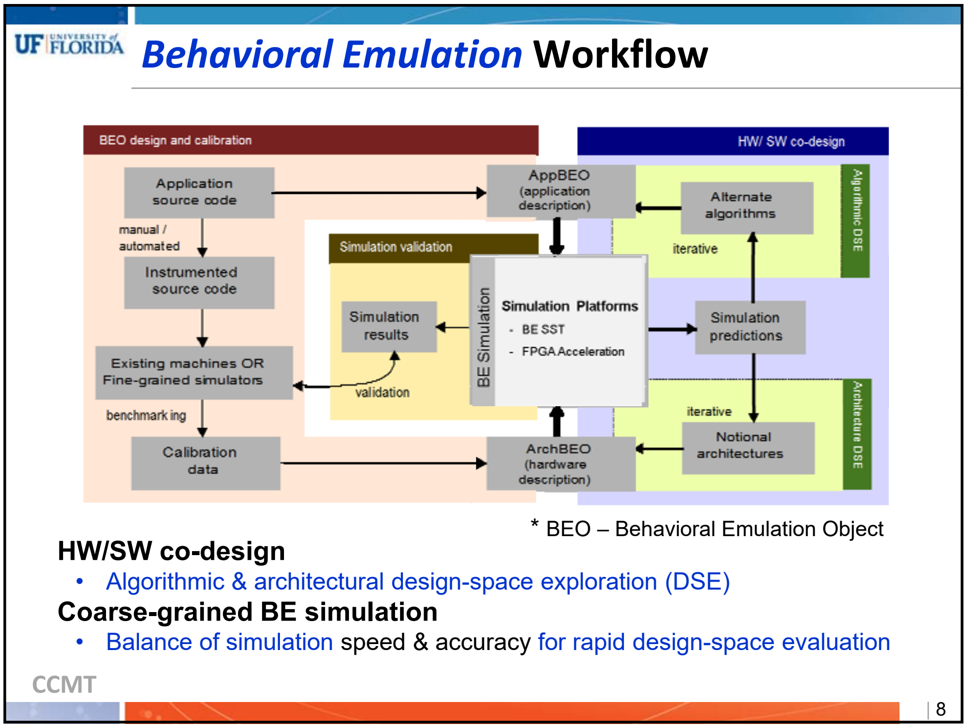

Behavioral Emulation Workflow

* BEO – Behavioral Emulation Object

HW/SW co-design

- Algorithmic & architectural design-space exploration (DSE)

Coarse-grained BE simulation

- Balance of simulation speed & accuracy for rapid design-space evaluation

CCMT

BE simulations of HPC systems

- Titan @ ORNL
  - Cray XK7 architecture with Cray Gemini interconnect
  - 16-core AMD opterons; 18k nodes; 300k cores
  - 18k K20X Kepler GPUs
  - 32GB + 6GB memory/node

- Vulcan @ LLNL
  - IBM BG/Q architecture
  - 16 cores/node, 24k nodes, 390k cores
  - 16GB memory/node

- Quartz @ LLNL
  - Intel Xeon (Broadwell) with Intel Omni-Path interconnect
  - 36-core Intel Xeon E5-2695; 2,634 nodes; 94k cores
  - 128GB memory/node

Application Case Study: CMT-bone-BE

- CMT-nek & CMT-bone (mini-app) both are large codes
- To support extensive DSE*, we need:
  - Key compute kernels & comm. patterns that affect performance
  - **Abstract, modular,** easy to **modify** and **instrument** for algorithmic DSE

CMT-nek Workflow
(highlighted portions represented in CMT-bone-BE)

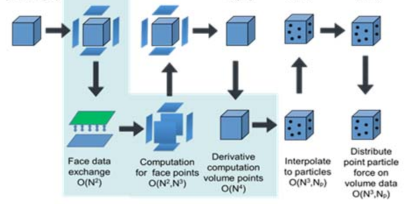

- Computation- & communication-intensive portions of CMT-nek workflow in C & MPI
  - Volume-to-surface data extraction
  - Face data exchange (with neighbors)
  - Derivative computation volume points
- Easier to modify for algorithmic DSE
- 1000 vs 10s of thousands lines of code

* design-space exploration

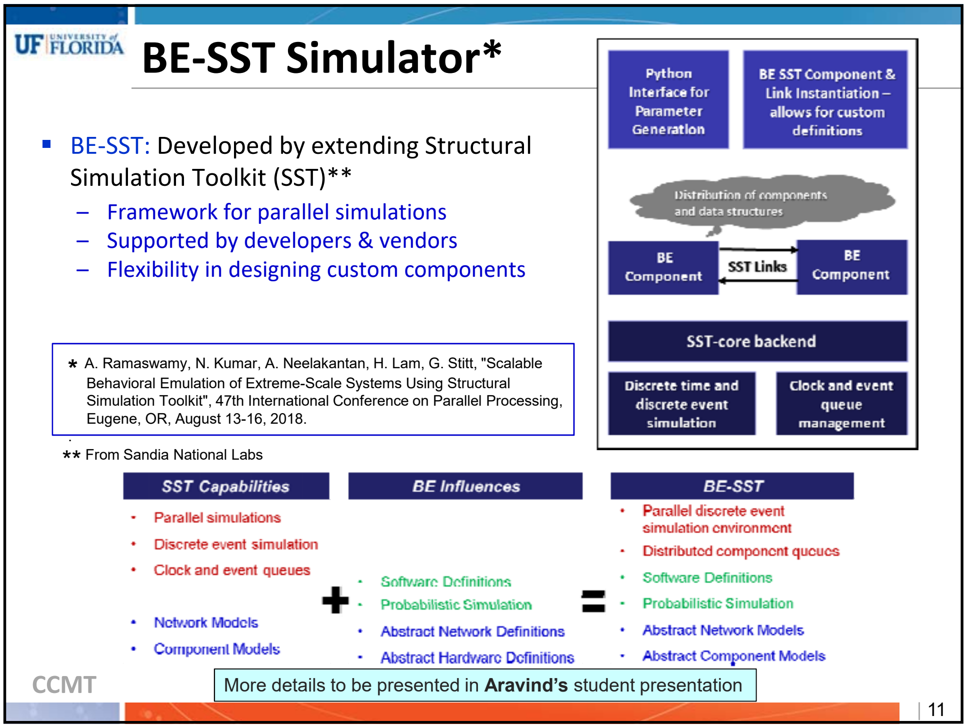

BE-SST Simulator*

- BE-SST: Developed by extending Structural Simulation Toolkit (SST)**
  - Framework for parallel simulations
  - Supported by developers & vendors
  - Flexibility in designing custom components

Python Interface for Parameter Generation

BE SST Component & Link Instantiation – allows for custom definitions

Distribution of components and data structures

BE Component

SST Links

BE Component

SST-core backend

Discrete time and discrete event simulation

Clock and event queue management

* A. Ramaswamy, N. Kumar, A. Neelakantan, H. Lam, G. Stitt, "Scalable Behavioral Emulation of Extreme-Scale Systems Using Structural Simulation Toolkit", 47th International Conference on Parallel Processing, Eugene, OR, August 13-16, 2018.

** From Sandia National Labs

SST Capabilities
- Parallel simulations
- Discrete event simulation
- Clock and event queues
- Network Models
- Component Models

\+

BE Influences
- Software Definitions
- Probabilistic Simulation
- Abstract Network Definitions
- Abstract Hardware Definitions

=

BE-SST
- Parallel discrete event simulation environment
- Distributed component queues
- Software Definitions
- Probabilistic Simulation
- Abstract Network Models
- Abstract Component Models

More details to be presented in **Aravind's** student presentation

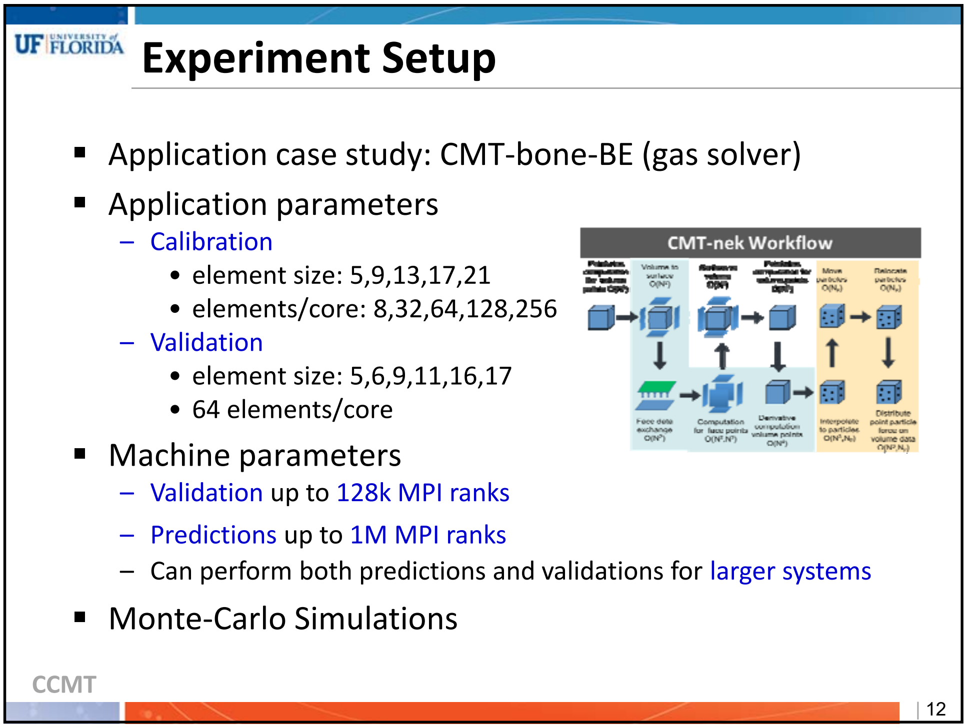

Experiment Setup

- Application case study: CMT-bone-BE (gas solver)
- Application parameters
  - Calibration
    - element size: 5,9,13,17,21
    - elements/core: 8,32,64,128,256
  - Validation
    - element size: 5,6,9,11,16,17
    - 64 elements/core
- Machine parameters
  - Validation up to 128k MPI ranks
  - Predictions up to 1M MPI ranks
  - Can perform both predictions and validations for larger systems
- Monte-Carlo Simulations

CMT-nek Workflow

UF FLORIDA

Recap: BE Simulations of Titan+ & Vulcan+

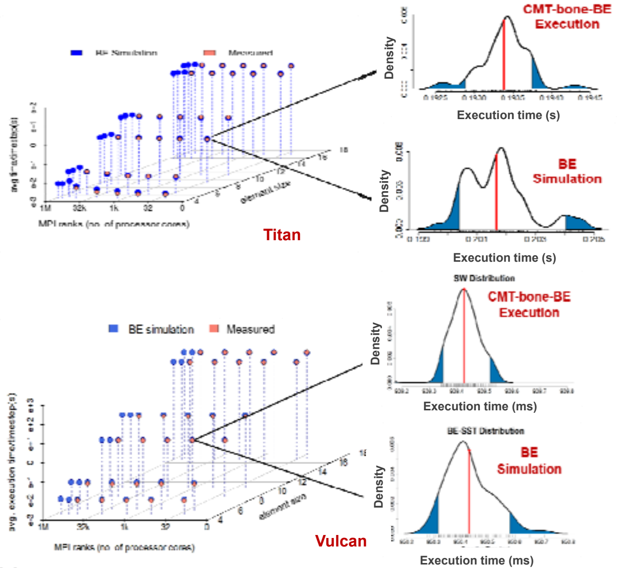

- Simulating bigger system than Titan and Vulcan (approaching 1 Million cores)
- Average % error between CMT-bone-BE simulation and execution time is 4%

CCMT

More details to be presented in **Aravind's** student presentation

---

UF FLORIDA

Architectural DSE

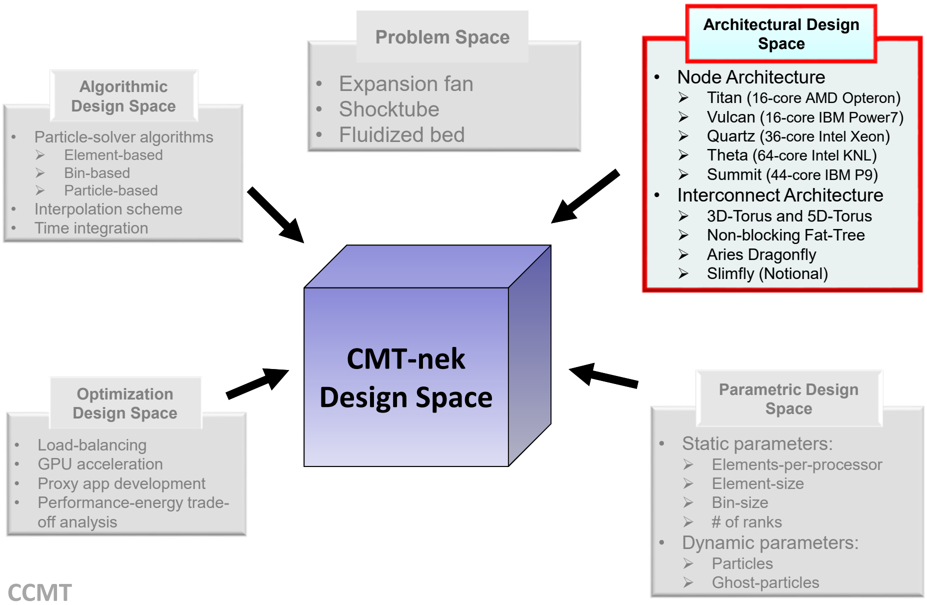

Algorithmic Design Space
- Particle-solver algorithms
  - Element-based
  - Bin-based
  - Particle-based
- Interpolation scheme
- Time integration

Problem Space
- Expansion fan
- Shocktube
- Fluidized bed

Architectural Design Space
- Node Architecture
  - Titan (16-core AMD Opteron)
  - Vulcan (16-core IBM Power7)
  - Quartz (36-core Intel Xeon)
  - Theta (64-core Intel KNL)
  - Summit (44-core IBM P9)
- Interconnect Architecture
  - 3D-Torus and 5D-Torus
  - Non-blocking Fat-Tree
  - Aries Dragonfly
  - Slimfly (Notional)

CMT-nek Design Space

Optimization Design Space
- Load-balancing
- GPU acceleration
- Proxy app development
- Performance-energy trade-off analysis

Parametric Design Space
- Static parameters:
  - Elements-per-processor
  - Element-size
  - Bin-size
  - # of ranks
- Dynamic parameters:
  - Particles
  - Ghost-particles

CCMT

Architectural DSE

- BE-SST performance validated in the past
- Currently, moving towards notional architecture
- Plug-n-play architectural DSE
  - Mixing node architecture with interconnect architecture

Architectural Design Space

- Node Architecture
  - Titan (16-core AMD Opteron)
  - Vulcan (16-core IBM Power7)
  - Quartz (36-core Intel Xeon)
  - Theta (64-core Intel KNL)
  - Summit (44-core IBM P9)
- Interconnect Architecture
  - 3D-Torus and 5D-Torus
  - Non-blocking Fat-Tree
  - Aries Dragonfly
  - Slimfly (Notional)

Machine	Node architecture	Interconnect
1	36-core Intel Xeon (Quartz architecture)	5D torus (of Vulcan)
2	36-core Intel Xeon (Quartz architecture)	3D torus (of Titan)

BE Simulations: Node Architecture

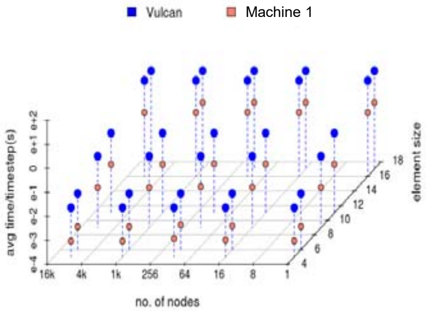

Vulcan vs Quartz with Vulcan interconnect

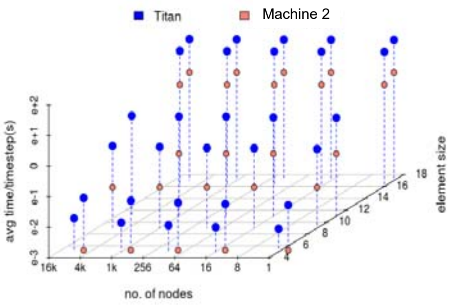

Titan vs Quartz with Titan interconnect

- Comparing compute node performance of Vulcan and Titan with Quartz
- Machine 1 and Machine 2 (with Quartz compute node) is faster than Vulcan compute nodes (by 21x) and Titan compute node (by 5x)
  - Expected behavior – Quartz's Intel Xeon is newer and faster than Vulcan's Power7 and Titan's AMD Opteron

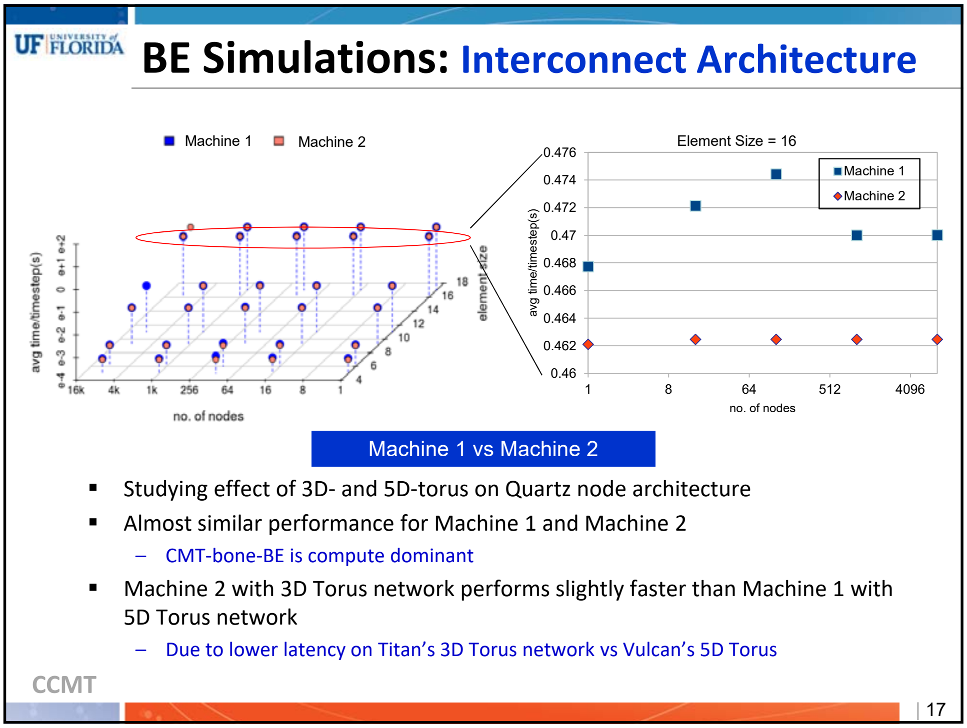

BE Simulations: Interconnect Architecture

Machine 1 vs Machine 2

- Studying effect of 3D- and 5D-torus on Quartz node architecture
- Almost similar performance for Machine 1 and Machine 2
  - CMT-bone-BE is compute dominant
- Machine 2 with 3D Torus network performs slightly faster than Machine 1 with 5D Torus network
  - Due to lower latency on Titan's 3D Torus network vs Vulcan's 5D Torus

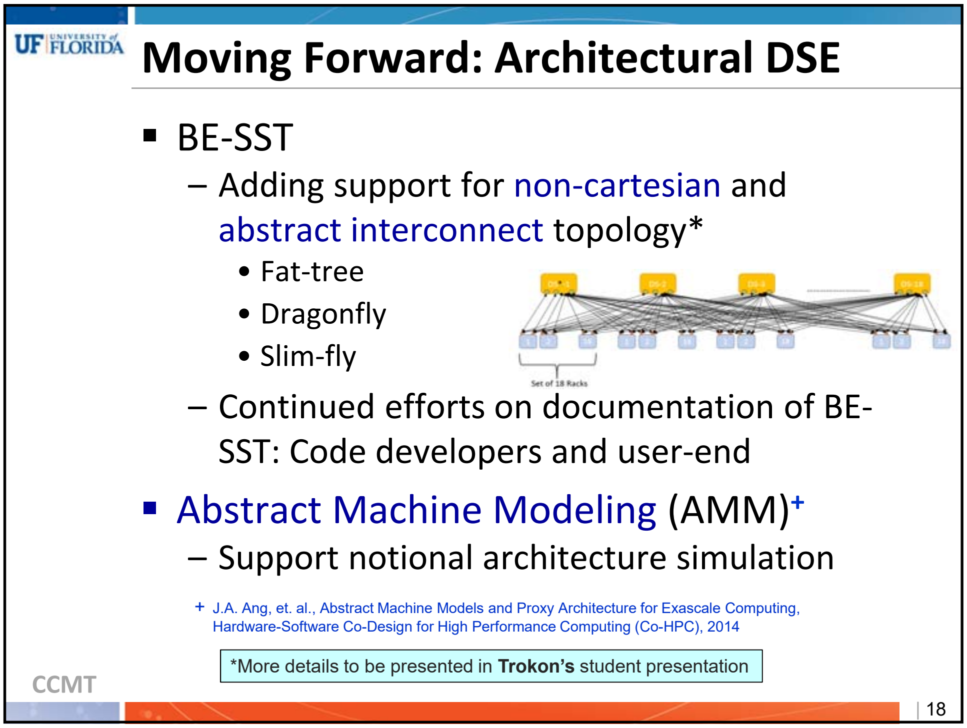

Moving Forward: Architectural DSE

- BE-SST
  - Adding support for non-cartesian and abstract interconnect topology*
    - Fat-tree
    - Dragonfly
    - Slim-fly
  - Continued efforts on documentation of BE-SST: Code developers and user-end
- Abstract Machine Modeling (AMM)+
  - Support notional architecture simulation

+ J.A. Ang, et. al., Abstract Machine Models and Proxy Architecture for Exascale Computing, Hardware-Software Co-Design for High Performance Computing (Co-HPC), 2014

*More details to be presented in **Trokon's** student presentation

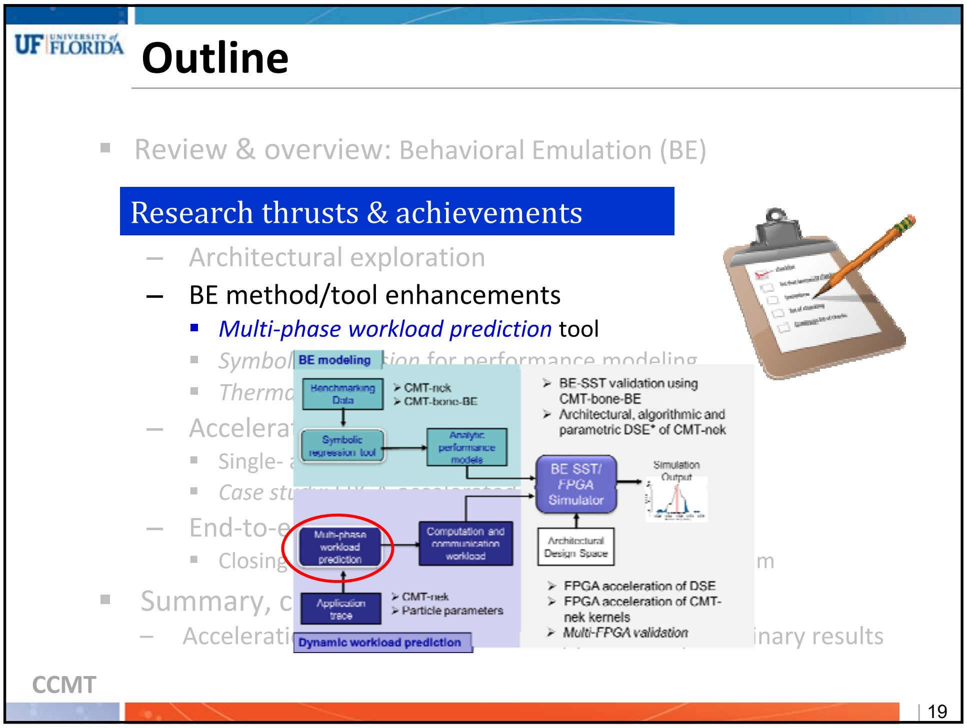

Outline

- Review & overview: Behavioral Emulation (BE)

Research thrusts & achievements

- Architectural exploration
- BE method/tool enhancements
  - *Multi-phase workload prediction* tool
  - *Symbol* *ion for performance modeling*
  - *Thermo*
- Accelera
  - Single- a
  - *Case stu*
- End-to-e
  - Closing m
- Summary, c
  - Accelerati inary results

CCMT

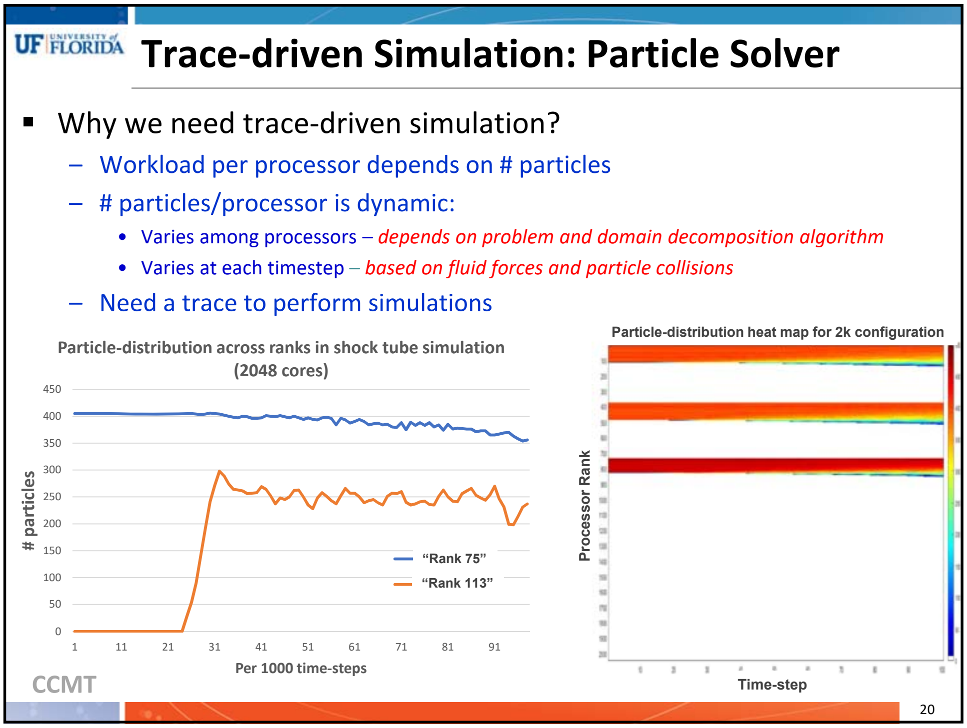

Trace-driven Simulation: Particle Solver

- Why we need trace-driven simulation?
  - Workload per processor depends on # particles
  - # particles/processor is dynamic:
    - Varies among processors – *depends on problem and domain decomposition algorithm*
    - Varies at each timestep – *based on fluid forces and particle collisions*
  - Need a trace to perform simulations

Particle-distribution across ranks in shock tube simulation (2048 cores)

Particle-distribution heat map for 2k configuration

CCMT

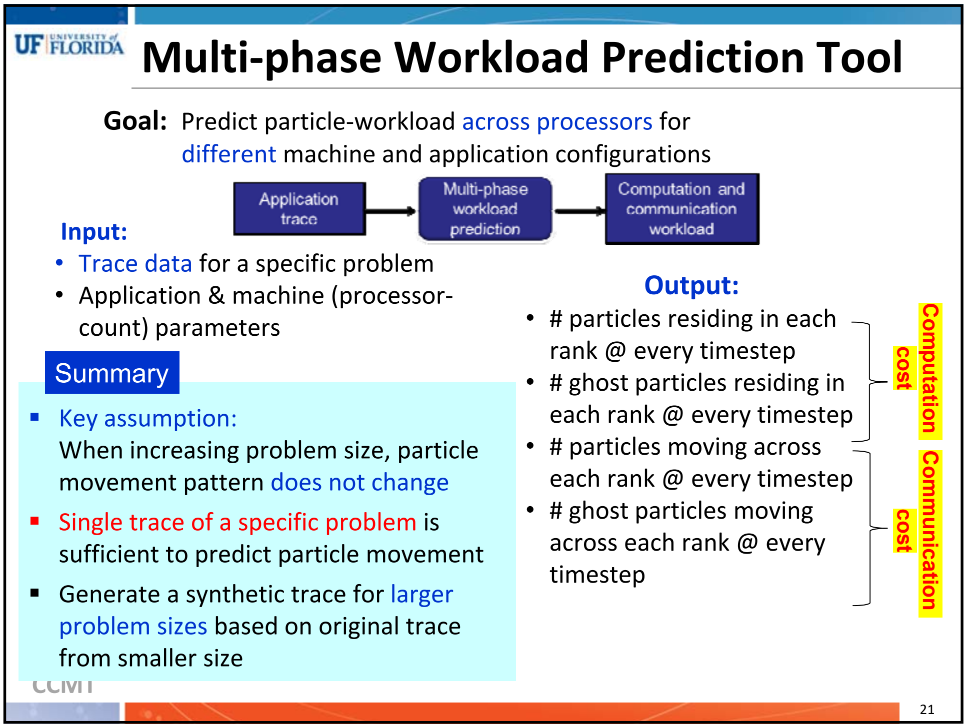

Multi-phase Workload Prediction Tool

Goal: Predict particle-workload across processors for different machine and application configurations

Input:

- Trace data for a specific problem
- Application & machine (processor-count) parameters

Summary

- Key assumption:
  When increasing problem size, particle movement pattern does not change
- Single trace of a specific problem is sufficient to predict particle movement
- Generate a synthetic trace for larger problem sizes based on original trace from smaller size

Output:

- # particles residing in each rank @ every timestep
- # ghost particles residing in each rank @ every timestep
- # particles moving across each rank @ every timestep
- # ghost particles moving across each rank @ every timestep

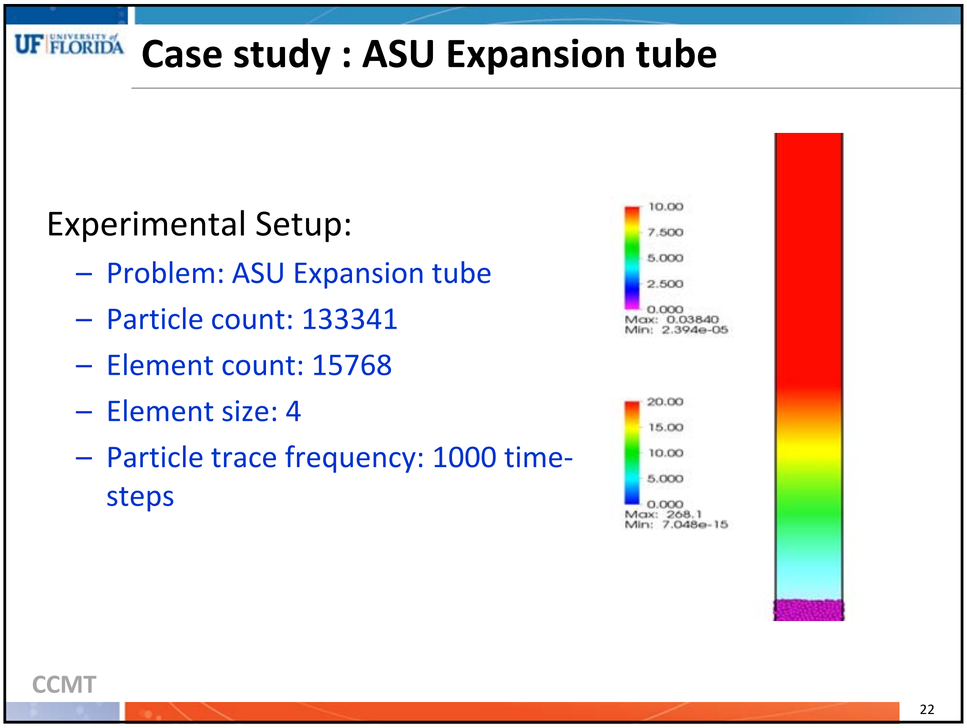

Case study : ASU Expansion tube

Experimental Setup:

- Problem: ASU Expansion tube
- Particle count: 133341
- Element count: 15768
- Element size: 4
- Particle trace frequency: 1000 time-steps

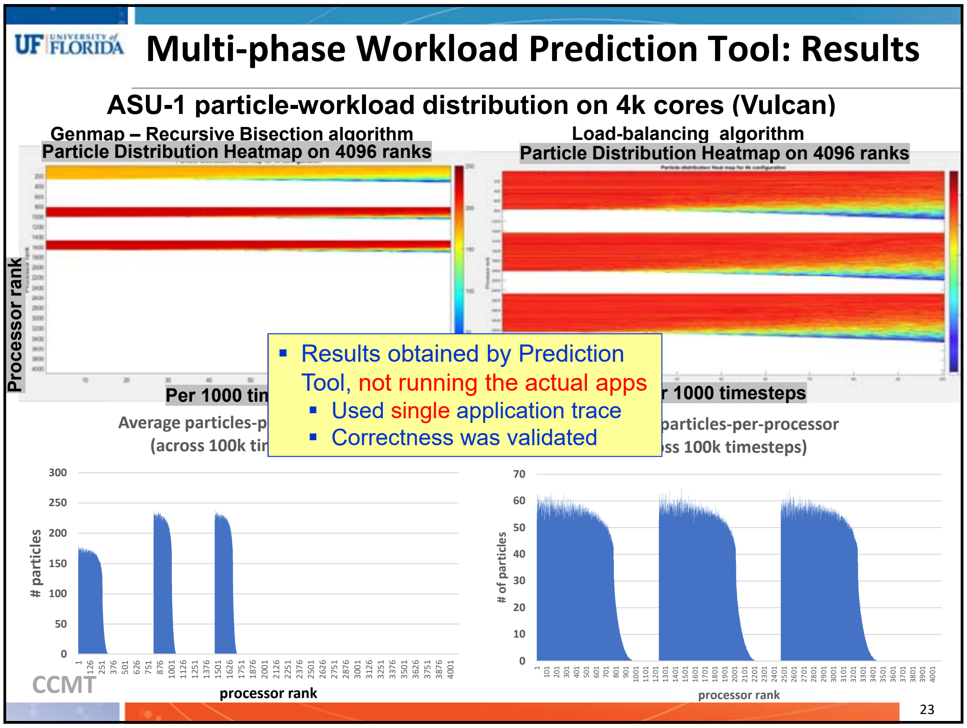

Multi-phase Workload Prediction Tool: Results

ASU-1 particle-workload distribution on 4k cores (Vulcan)

Genmap – Recursive Bisection algorithm

Particle Distribution Heatmap on 4096 ranks

Load-balancing algorithm

Particle Distribution Heatmap on 4096 ranks

- Results obtained by Prediction Tool, not running the actual apps
  - Used single application trace
  - Correctness was validated

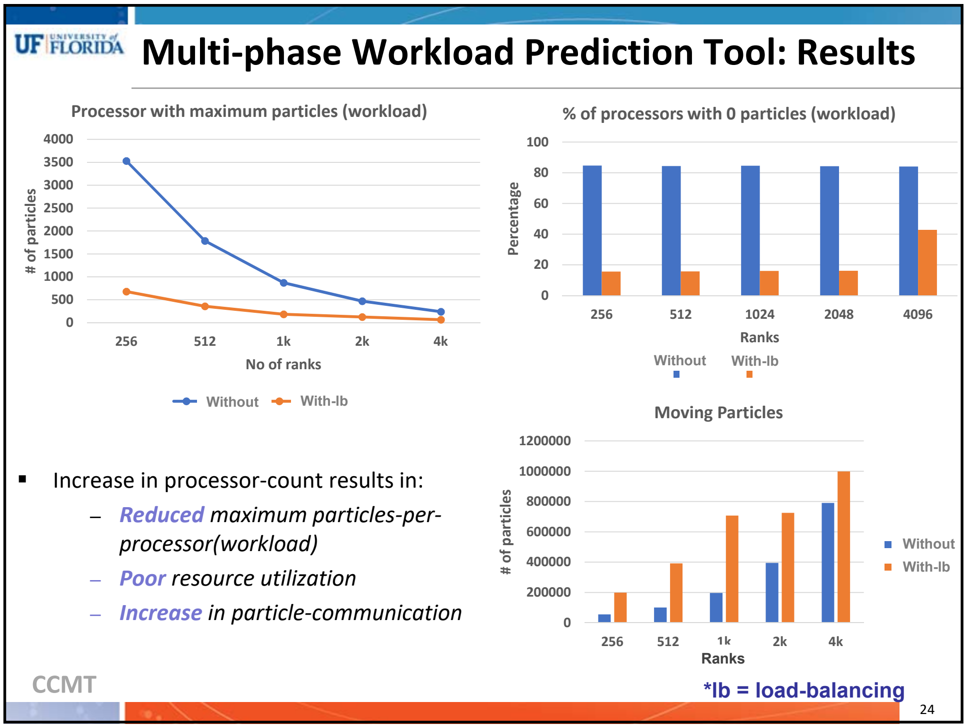

Multi-phase Workload Prediction Tool: Results

- Increase in processor-count results in:
  - *Reduced maximum particles-per-processor(workload)*
  - *Poor resource utilization*
  - *Increase in particle-communication*

***lb = load-balancing**

CCMT

Outline

- Review & overview: Behavioral Emulation (BE)

Research thrusts & achievements

  - Architectural exploration
  - BE method/tool enhancements
    - *Particle-workload* prediction tool
    - *Symbolic regression* for performance modeling
    - *Therm* ments & modeling
  - Acceler
    - Single
    - *Case s*
  - End-to-
    - Closin eam
- Summary,
  - Accelerat ninary results

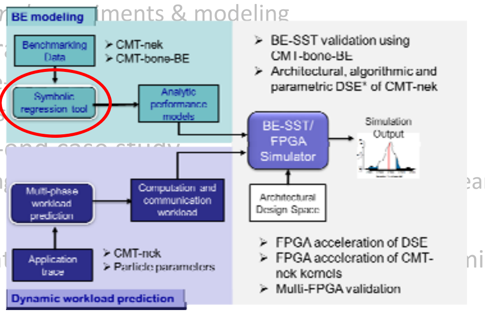

Symbolic Regression: Motivation

- Motivation:
  - Multi-parameter performance modeling is challenging
  - Requires accuracy, computational efficiency, simple implementation, flexibility
- Our previous approaches:
  - N-dimensional interpolation:
    - Computationally expensive
    - Difficult to implement, needs customization for each example
    - Insufficient accuracy
  - Multiple linear regression:
    - Faster than interpolation, relatively simple to implement
    - Need thorough understanding of code to generate performance models
    - May not capture machine-specific behaviors that deviate from basic model

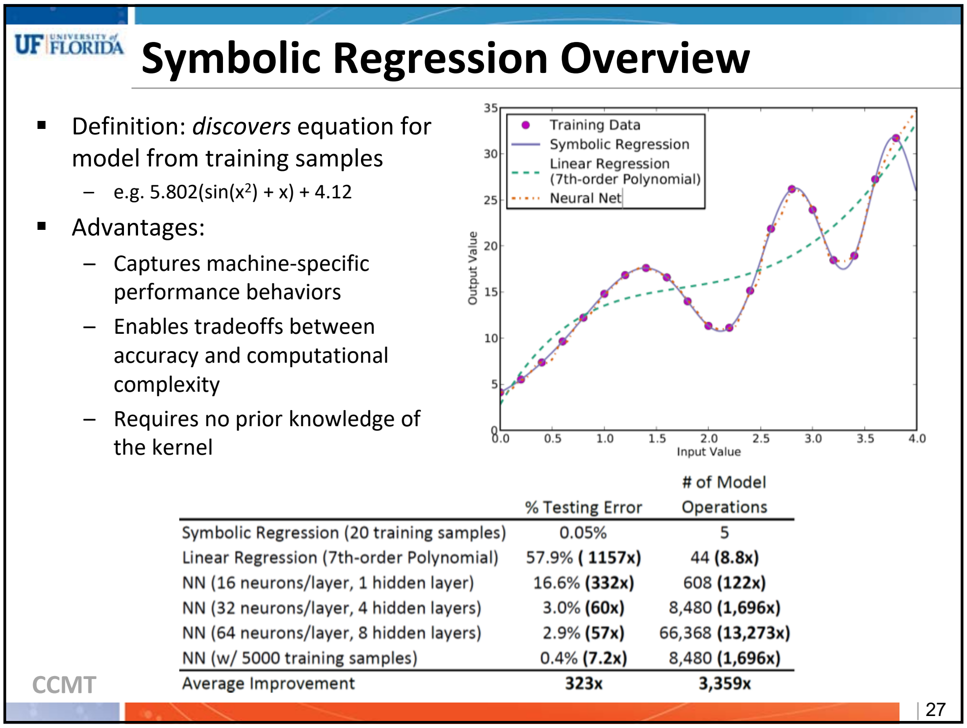

Symbolic Regression Overview

- Definition: *discovers* equation for model from training samples
  - e.g. 5.802(sin(x^2) + x) + 4.12
- Advantages:
  - Captures machine-specific performance behaviors
  - Enables tradeoffs between accuracy and computational complexity
  - Requires no prior knowledge of the kernel

	% Testing Error	# of Model Operations
Symbolic Regression (20 training samples)	0.05%	5
Linear Regression (7th-order Polynomial)	57.9% **(1157x)**	44 **(8.8x)**
NN (16 neurons/layer, 1 hidden layer)	16.6% **(332x)**	608 **(122x)**
NN (32 neurons/layer, 4 hidden layers)	3.0% **(60x)**	8,480 **(1,696x)**
NN (64 neurons/layer, 8 hidden layers)	2.9% **(57x)**	66,368 **(13,273x)**
NN (w/ 5000 training samples)	0.4% **(7.2x)**	8,480 **(1,696x)**
Average Improvement	**323x**	**3,359x**

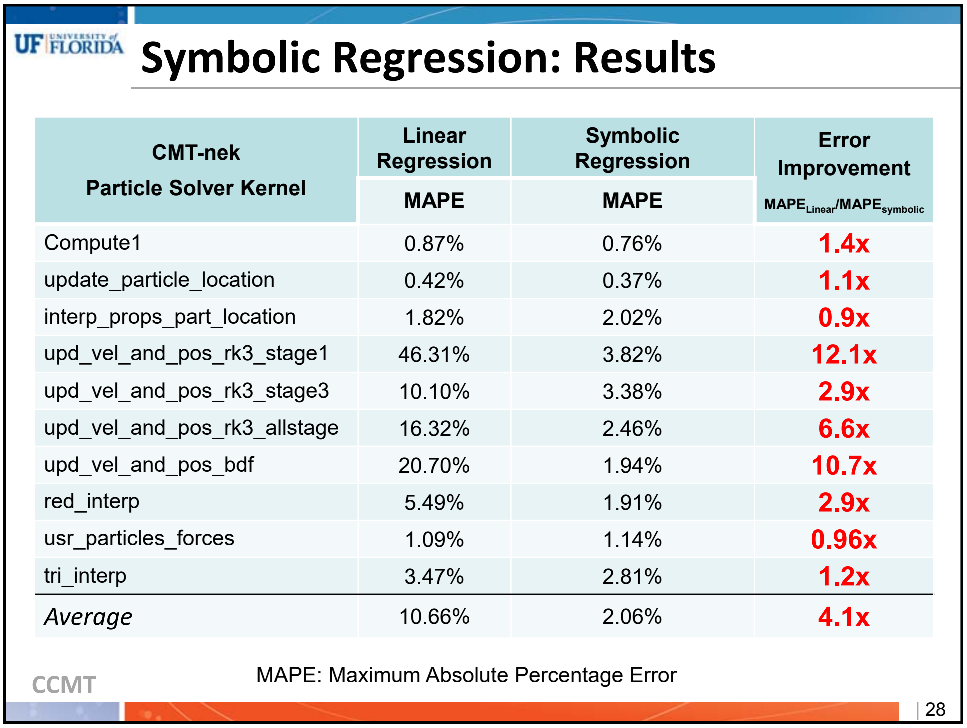

Symbolic Regression: Results

CMT-nek Particle Solver Kernel	Linear Regression	Symbolic Regression	Error Improvement
	MAPE	MAPE	$MAPE_{Linear}/MAPE_{symbolic}$
Compute1	0.87%	0.76%	**1.4x**
update_particle_location	0.42%	0.37%	**1.1x**
interp_props_part_location	1.82%	2.02%	**0.9x**
upd_vel_and_pos_rk3_stage1	46.31%	3.82%	**12.1x**
upd_vel_and_pos_rk3_stage3	10.10%	3.38%	**2.9x**
upd_vel_and_pos_rk3_allstage	16.32%	2.46%	**6.6x**
upd_vel_and_pos_bdf	20.70%	1.94%	**10.7x**
red_interp	5.49%	1.91%	**2.9x**
usr_particles_forces	1.09%	1.14%	**0.96x**
tri_interp	3.47%	2.81%	**1.2x**
Average	10.66%	2.06%	**4.1x**

MAPE: Maximum Absolute Percentage Error

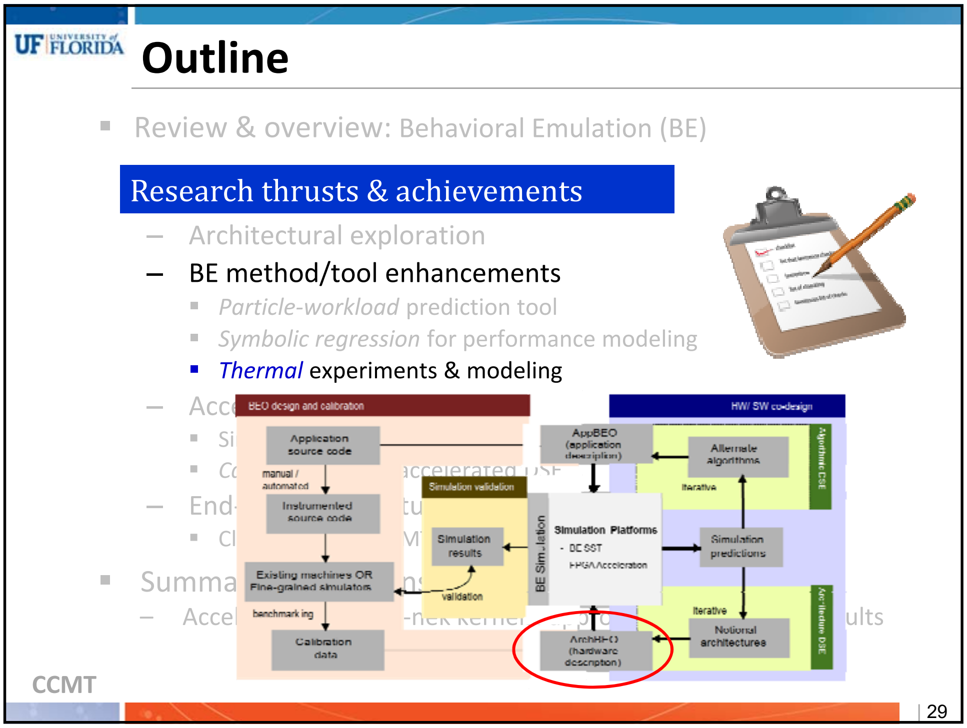

Outline

- Review & overview: Behavioral Emulation (BE)

Research thrusts & achievements

- Architectural exploration
- BE method/tool enhancements
  - *Particle-workload* prediction tool
  - *Symbolic regression* for performance modeling
  - *Thermal* experiments & modeling
- Acc
  - Si
  - Cc accelerated DSE
- End- tu
  - Cl MT
- Summa ns
  - Accel -nek kernel app ults

CCMT

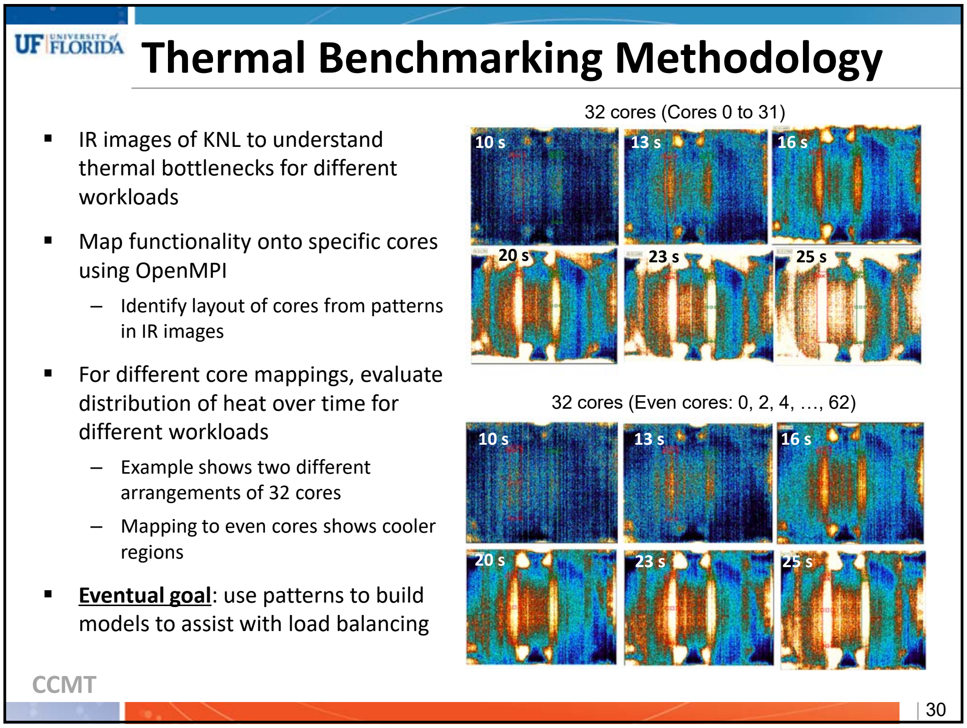

Thermal Benchmarking Methodology

- IR images of KNL to understand thermal bottlenecks for different workloads
- Map functionality onto specific cores using OpenMPI
  - Identify layout of cores from patterns in IR images
- For different core mappings, evaluate distribution of heat over time for different workloads
  - Example shows two different arrangements of 32 cores
  - Mapping to even cores shows cooler regions
- **Eventual goal**: use patterns to build models to assist with load balancing

CCMT

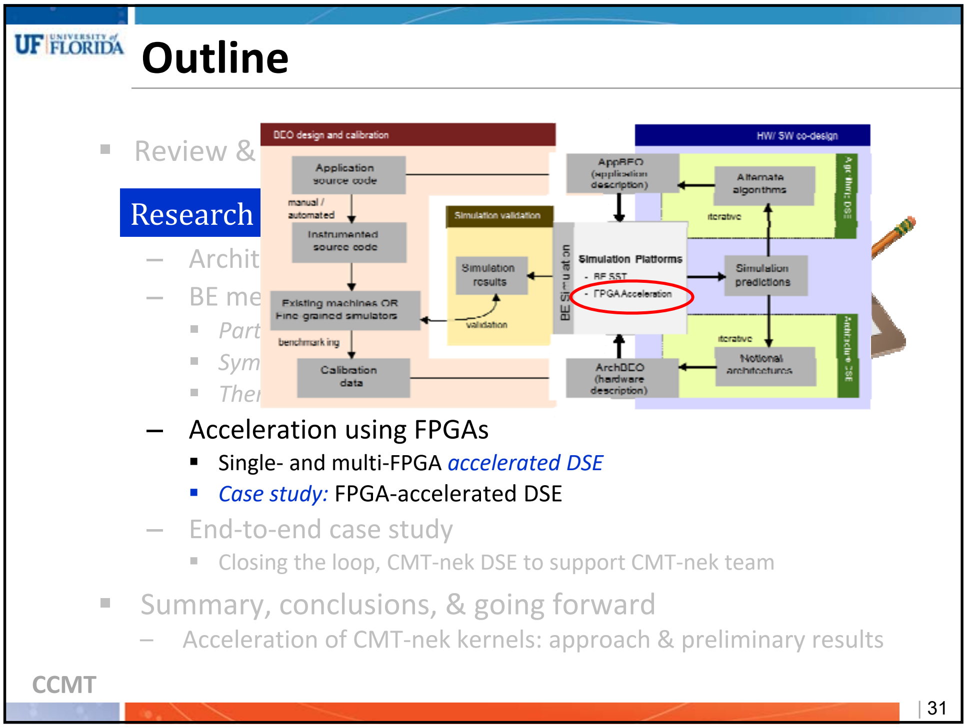

Outline

- Review &
- Research
  - Archit
  - BE me
    - Part
    - Sym
    - Ther

  - Acceleration using FPGAs
    - Single- and multi-FPGA *accelerated DSE*
    - *Case study:* FPGA-accelerated DSE
  - End-to-end case study
    - Closing the loop, CMT-nek DSE to support CMT-nek team
- Summary, conclusions, & going forward
  - Acceleration of CMT-nek kernels: approach & preliminary results

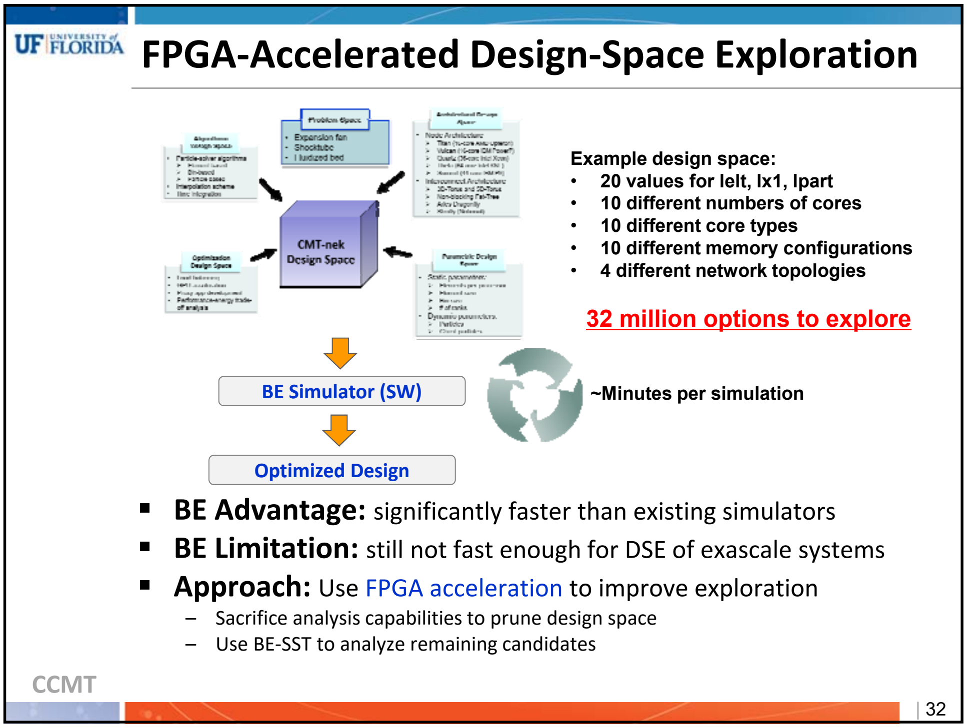

FPGA-Accelerated Design-Space Exploration

Example design space:

- **20 values for lelt, lx1, lpart**
- **10 different numbers of cores**
- **10 different core types**
- **10 different memory configurations**
- **4 different network topologies**

32 million options to explore

~Minutes per simulation

- **BE Advantage:** significantly faster than existing simulators
- **BE Limitation:** still not fast enough for DSE of exascale systems
- **Approach:** Use FPGA acceleration to improve exploration
  - Sacrifice analysis capabilities to prune design space
  - Use BE-SST to analyze remaining candidates

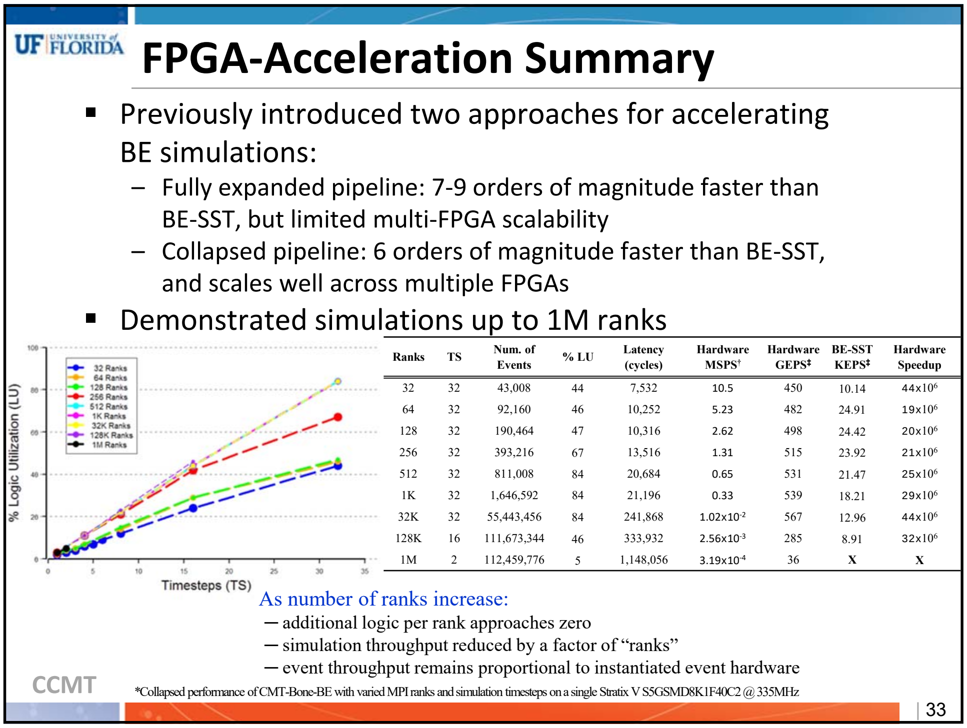

FPGA-Acceleration Summary

- Previously introduced two approaches for accelerating BE simulations:
  - Fully expanded pipeline: 7-9 orders of magnitude faster than BE-SST, but limited multi-FPGA scalability
  - Collapsed pipeline: 6 orders of magnitude faster than BE-SST, and scales well across multiple FPGAs
- Demonstrated simulations up to 1M ranks

Ranks	TS	Num. of Events	% LU	Latency (cycles)	Hardware MSPS†	Hardware GEPS‡	BE-SST KEPS‡	Hardware Speedup
32	32	43,008	44	7,532	10.5	450	10.14	44×10^6
64	32	92,160	46	10,252	5.23	482	24.91	19×10^6
128	32	190,464	47	10,316	2.62	498	24.42	20×10^6
256	32	393,216	67	13,516	1.31	515	23.92	21×10^6
512	32	811,008	84	20,684	0.65	531	21.47	25×10^6
1K	32	1,646,592	84	21,196	0.33	539	18.21	29×10^6
32K	32	55,443,456	84	241,868	1.02×10^{-2}	567	12.96	44×10^6
128K	16	111,673,344	46	333,932	2.56×10^{-3}	285	8.91	32×10^6
1M	2	112,459,776	5	1,148,056	3.19×10^{-4}	36	X	X

As number of ranks increase:
— additional logic per rank approaches zero
— simulation throughput reduced by a factor of "ranks"
— event throughput remains proportional to instantiated event hardware

*Collapsed performance of CMT-Bone-BE with varied MPI ranks and simulation timesteps on a single Stratix V S5GSMD8K1F40C2 @ 335MHz

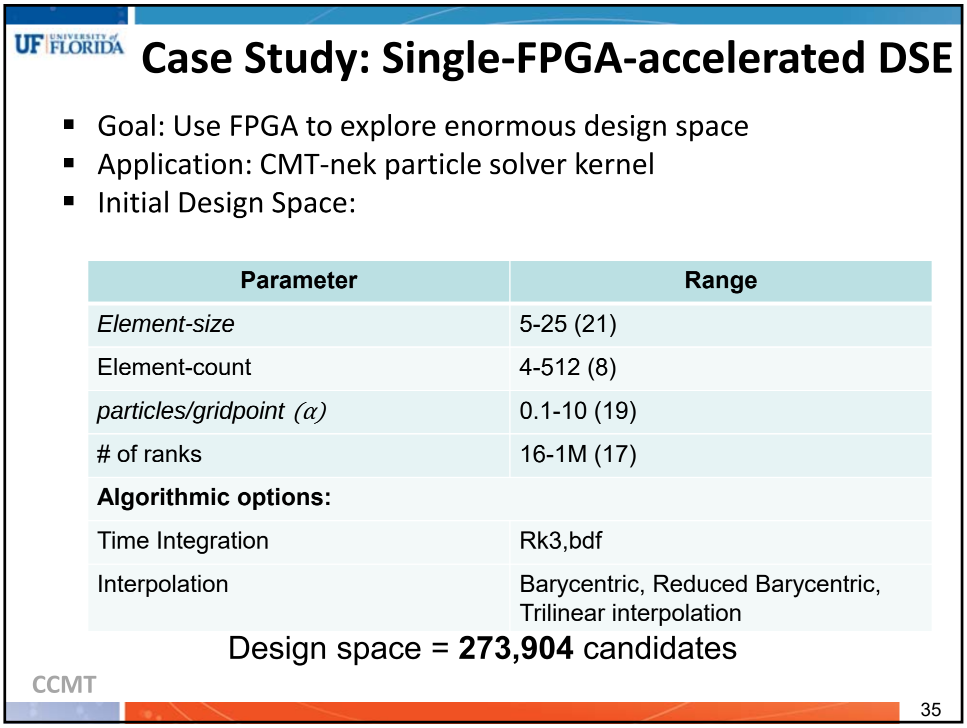

Case Study: Single-FPGA-accelerated DSE

- Goal: Use FPGA to explore enormous design space
- Application: CMT-nek particle solver kernel
- Initial Design Space:

Parameter	Range
Element-size	5-25 (21)
Element-count	4-512 (8)
particles/gridpoint (α)	0.1-10 (19)
# of ranks	16-1M (17)
Algorithmic options:	
Time Integration	Rk3,bdf
Interpolation	Barycentric, Reduced Barycentric, Trilinear interpolation

Design space = **273,904** candidates

Case Study: Single-FPGA-accelerated DSE

Ranks	No of events	FPGA Avg. Simulation time (microseconds)	BE-SST* Avg. Simulation time (seconds)	FPGA speedup
1024	24,576	3.53	8.234	2.33x10^6
2048	49,152	7.06	6.5	9.21x10^5
4096	98,304	14.12	17.3	1.22x10^6
8192	196,608	28.24	31	1.10x10^6
16k	393,216	56.48	120	2.12x10^6
32k	786,432	112.96	145	1.28x10^6
64k	1,572,864	225.92	405	1.79x10^6
128K	3,145,728	451.84	531	1.17x10^6
256K	6,291,456	903.68	1569	1.74x10^6
512K	12,582,912	1807.36	7416	4.10x10^6
1M	25,165,824	3614.72	X	X

- Each configuration (row) has **16,112** design candidates
- BE-SST takes **5.5 years** to explore the design space – ***FPGA takes 61 seconds***
  - **550 years** including 100 Monte Carlo simulations for each design candidate vs ***100 minutes***
  - **FPGA sacrifices analysis capabilities of BE-SST to achieve this speedup**

*ran on HiPerGator - 64 cores

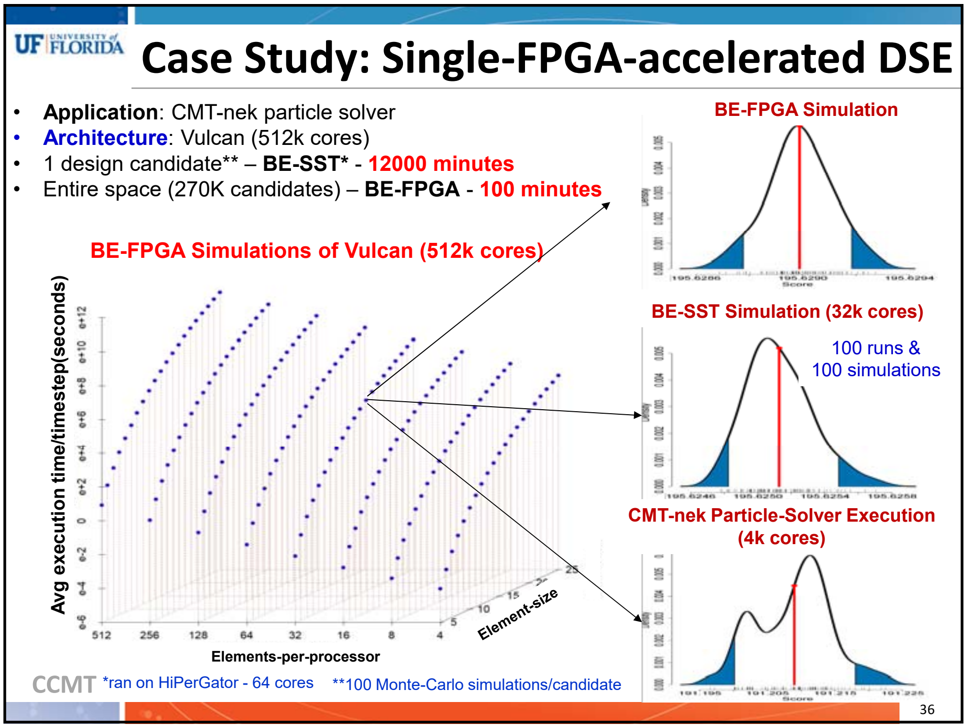

Case Study: Single-FPGA-accelerated DSE

- **Application**: CMT-nek particle solver
- **Architecture**: Vulcan (512k cores)
- 1 design candidate** – **BE-SST*** - **12000 minutes**
- Entire space (270K candidates) – **BE-FPGA** - **100 minutes**

BE-FPGA Simulations of Vulcan (512k cores)

BE-FPGA Simulation

BE-SST Simulation (32k cores)

100 runs & 100 simulations

CMT-nek Particle-Solver Execution (4k cores)

*ran on HiPerGator - 64 cores **100 Monte-Carlo simulations/candidate

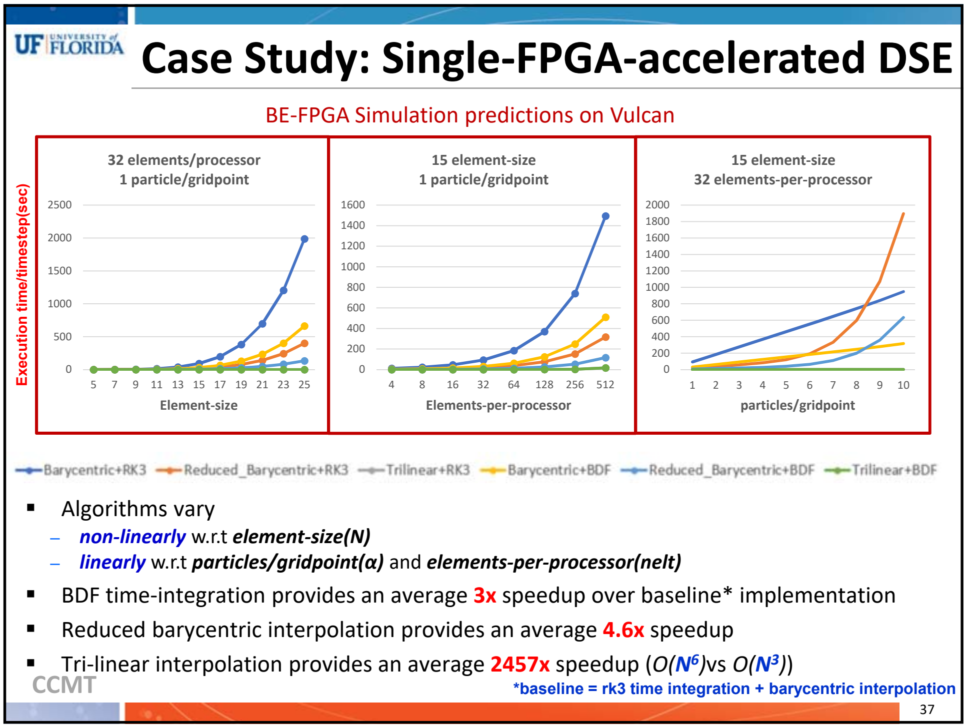

Case Study: Single-FPGA-accelerated DSE

BE-FPGA Simulation predictions on Vulcan

- Algorithms vary
  - ***non-linearly*** w.r.t ***element-size(N)***
  - ***linearly*** w.r.t ***particles/gridpoint(α)*** and ***elements-per-processor(nelt)***
- BDF time-integration provides an average **3x** speedup over baseline* implementation
- Reduced barycentric interpolation provides an average **4.6x** speedup
- Tri-linear interpolation provides an average **2457x** speedup (*O*(***N***^{6})vs *O*(***N***^{3}))

***baseline = rk3 time integration + barycentric interpolation**

37

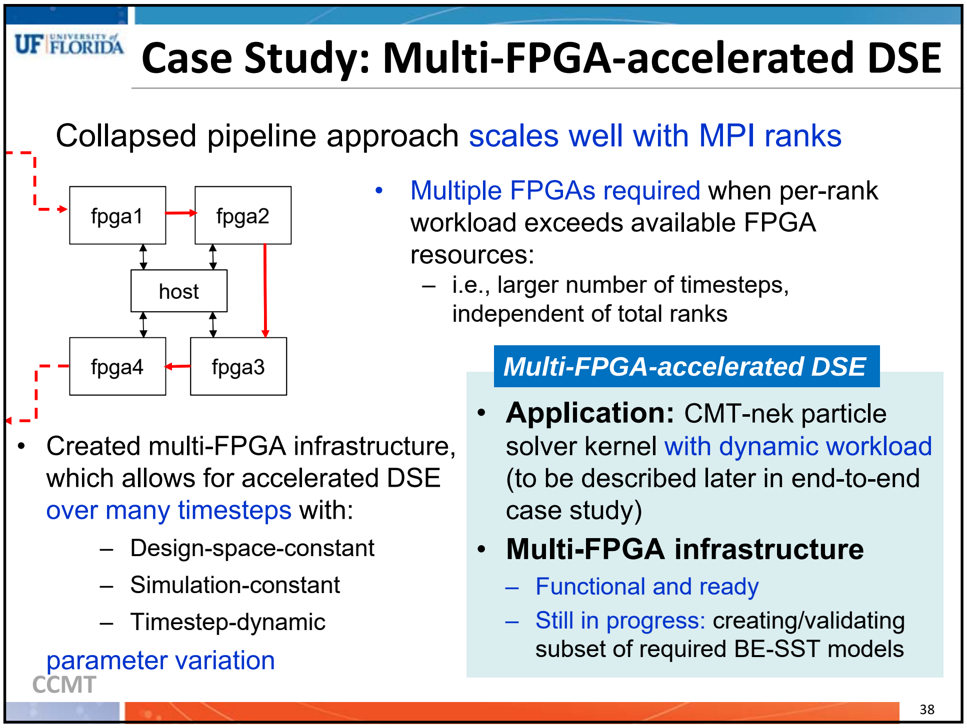

Case Study: Multi-FPGA-accelerated DSE

Collapsed pipeline approach scales well with MPI ranks

- Multiple FPGAs required when per-rank workload exceeds available FPGA resources:
  - i.e., larger number of timesteps, independent of total ranks

- Created multi-FPGA infrastructure, which allows for accelerated DSE over many timesteps with:
  - Design-space-constant
  - Simulation-constant
  - Timestep-dynamic

  parameter variation

Multi-FPGA-accelerated DSE

- **Application:** CMT-nek particle solver kernel with dynamic workload (to be described later in end-to-end case study)
- **Multi-FPGA infrastructure**
  - Functional and ready
  - Still in progress: creating/validating subset of required BE-SST models

38

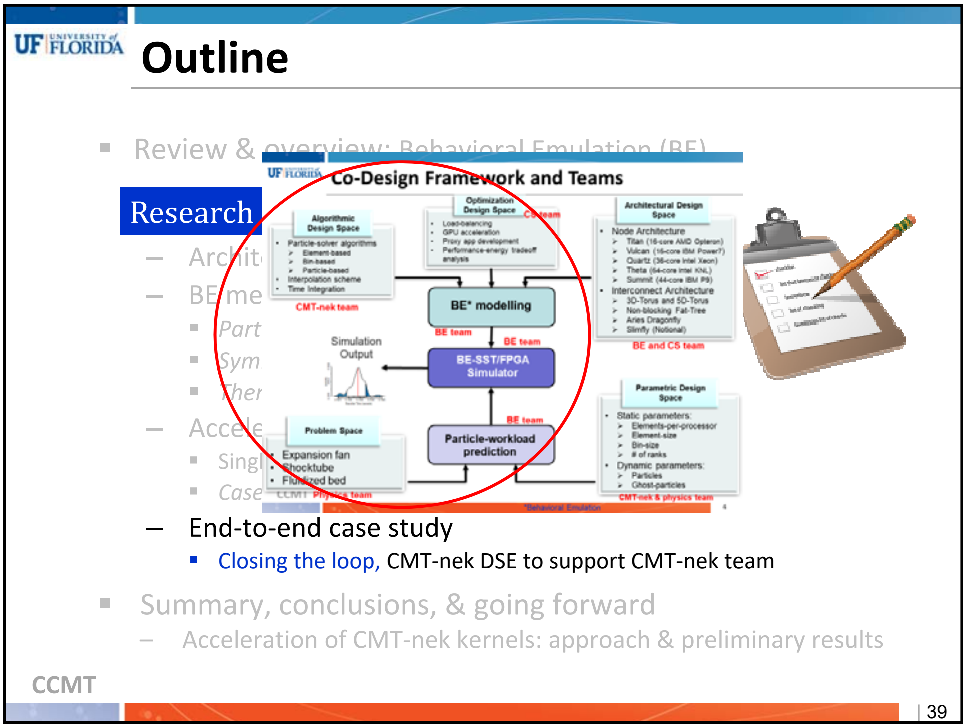

Outline

- Review & overview: Behavioral Emulation (BE)

Research
  - Archit
  - BE me
    - *Part*
    - *Sym*
    - *Ther*
  - Accele
    - Singl
    - *Case*
  - End-to-end case study
    - Closing the loop, CMT-nek DSE to support CMT-nek team
- Summary, conclusions, & going forward
  - Acceleration of CMT-nek kernels: approach & preliminary results

CCMT

39

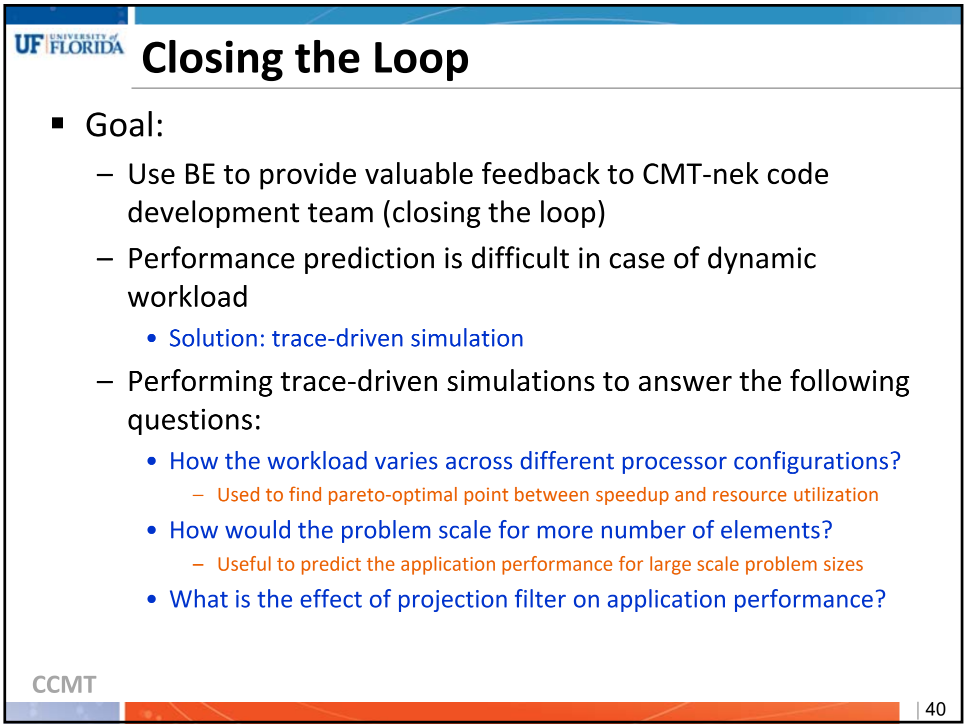

Closing the Loop

- Goal:
  - Use BE to provide valuable feedback to CMT-nek code development team (closing the loop)
  - Performance prediction is difficult in case of dynamic workload
    - Solution: trace-driven simulation
  - Performing trace-driven simulations to answer the following questions:
    - How the workload varies across different processor configurations?
      - Used to find pareto-optimal point between speedup and resource utilization
    - How would the problem scale for more number of elements?
      - Useful to predict the application performance for large scale problem sizes
    - What is the effect of projection filter on application performance?

CCMT

40

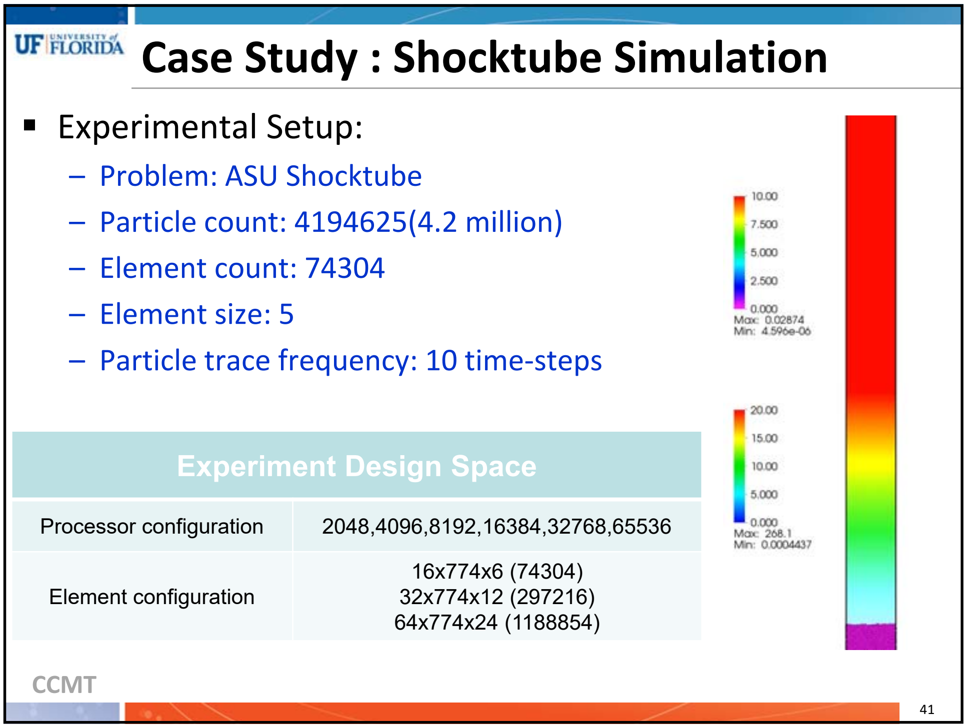

Case Study : Shocktube Simulation

- Experimental Setup:
  - Problem: ASU Shocktube
  - Particle count: 4194625(4.2 million)
  - Element count: 74304
  - Element size: 5
  - Particle trace frequency: 10 time-steps

Experiment Design Space	
Processor configuration	2048,4096,8192,16384,32768,65536
Element configuration	16x774x6 (74304) 32x774x12 (297216) 64x774x24 (1188854)

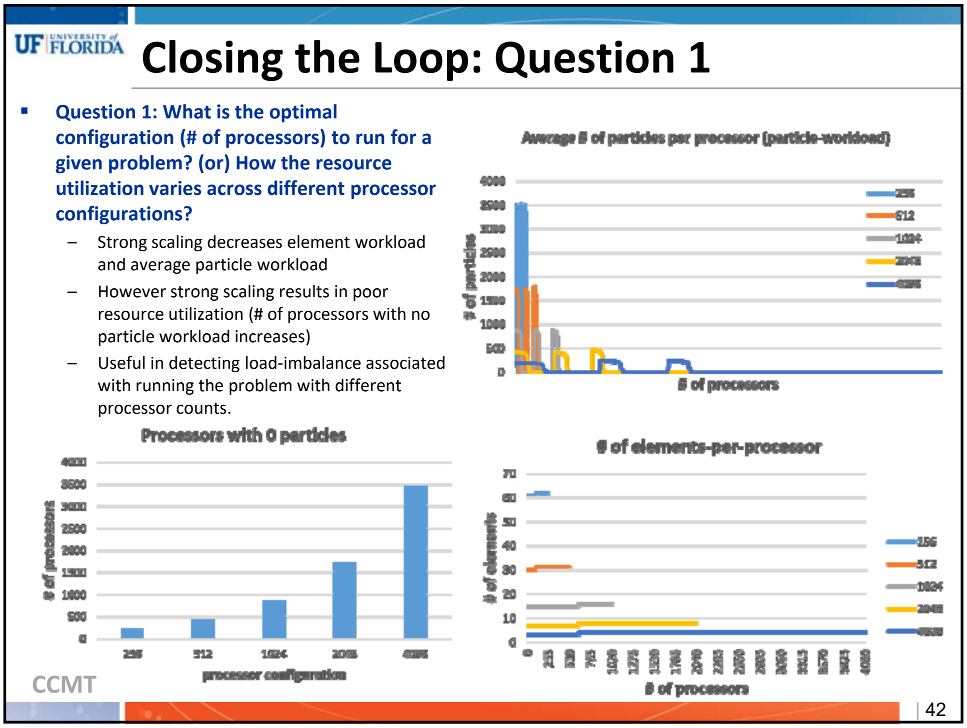

Closing the Loop: Question 1

- **Question 1: What is the optimal configuration (# of processors) to run for a given problem? (or) How the resource utilization varies across different processor configurations?**
  - Strong scaling decreases element workload and average particle workload
  - However strong scaling results in poor resource utilization (# of processors with no particle workload increases)
  - Useful in detecting load-imbalance associated with running the problem with different processor counts.

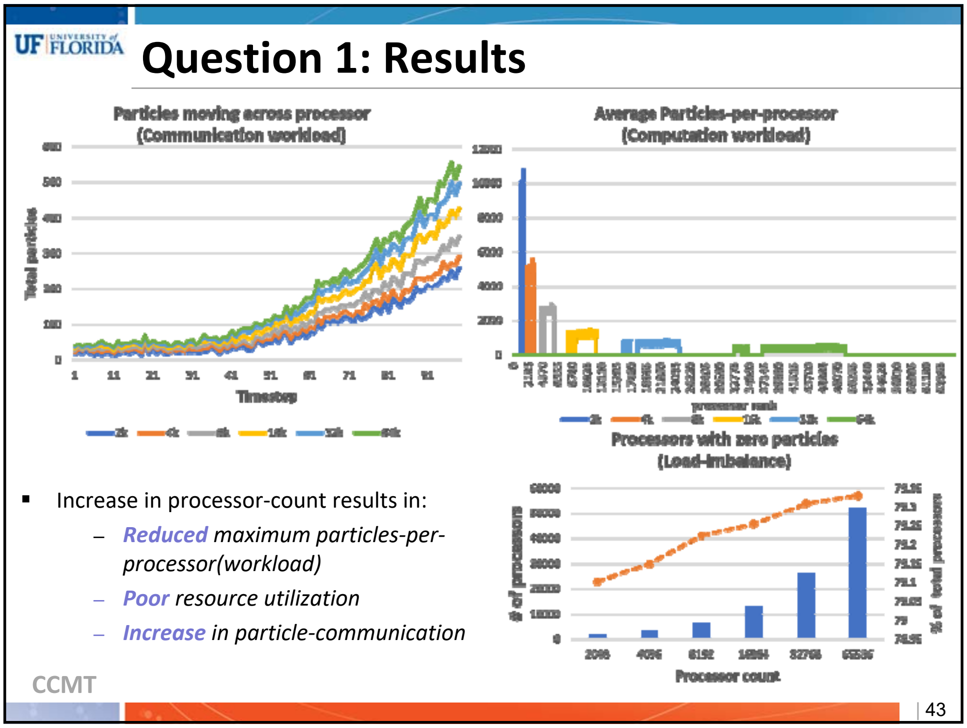

Question 1: Results

- Increase in processor-count results in:
  - *__Reduced__ maximum particles-per-processor(workload)*
  - *__Poor__ resource utilization*
  - *__Increase__ in particle-communication*

CCMT

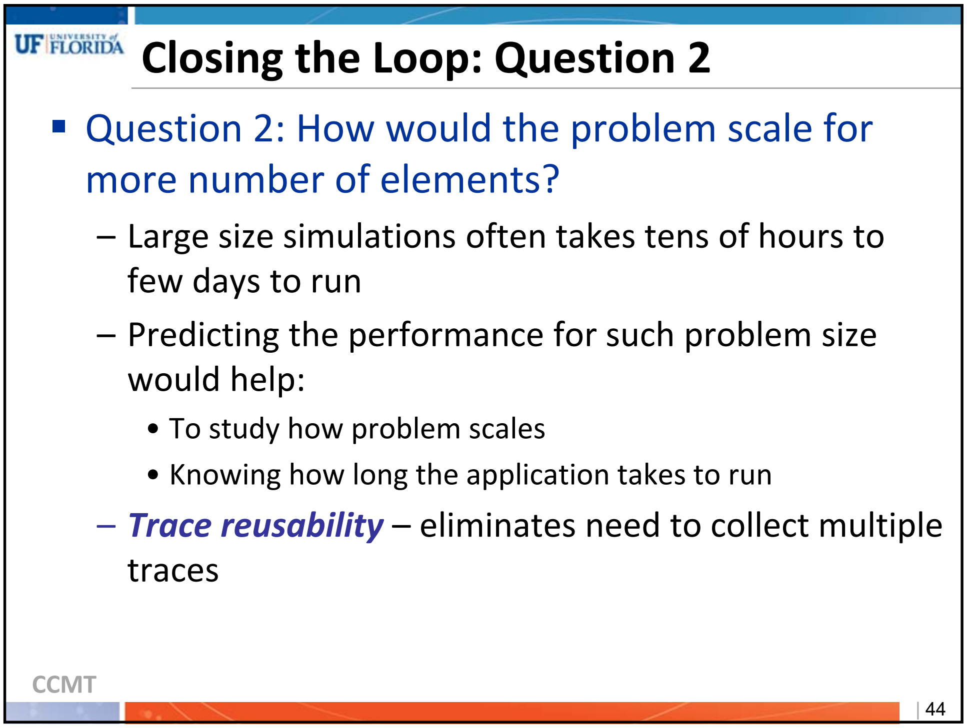

Closing the Loop: Question 2

- Question 2: How would the problem scale for more number of elements?
  - Large size simulations often takes tens of hours to few days to run
  - Predicting the performance for such problem size would help:
    - To study how problem scales
    - Knowing how long the application takes to run
  - ***Trace reusability*** – eliminates need to collect multiple traces

CCMT

UF FLORIDA

Question 2: Results

Prediction

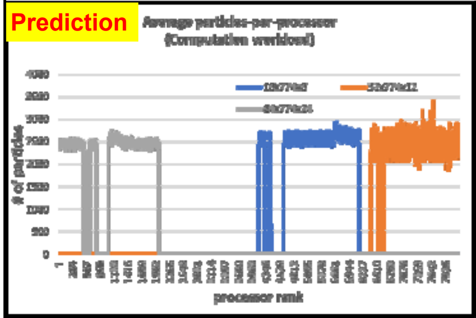

Validation

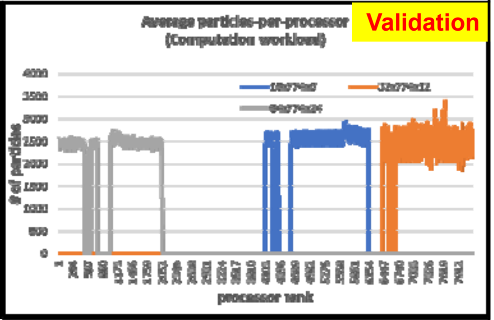

- Increase in ***element-count*** resulted in:
  - *No significant change* in particle-workload distribution
  - *Slight decrease* in load-imbalance
- Single trace was sufficient to predict accurately for larger element configurations

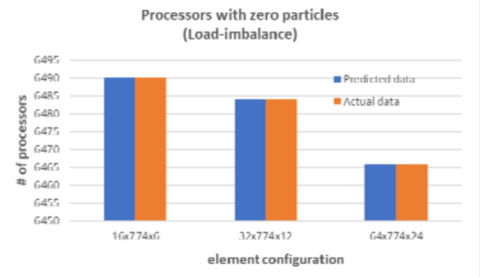

UF FLORIDA

Closing the Loop: Question 3

- **Question 3: What is the effect of projection filter on application performance?**
  - Ghost particles are generated based on projection filter-width
  - Different parameter values have a significant impact on the application performance
    - Larger bin-size - *less bins* but *more ghost particles*
    - Smaller bin-size – *more bins* but *less ghost particles*
- **Experiment in progress**

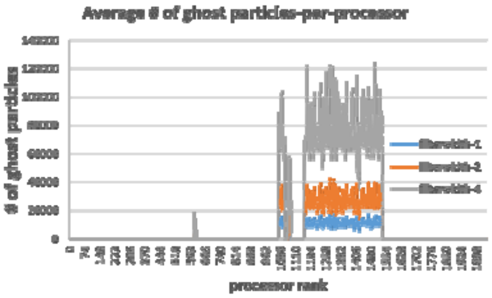

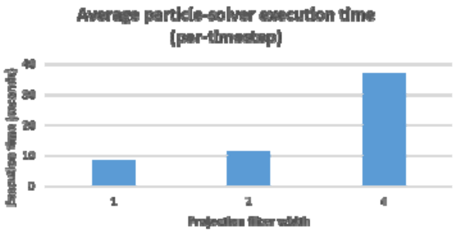

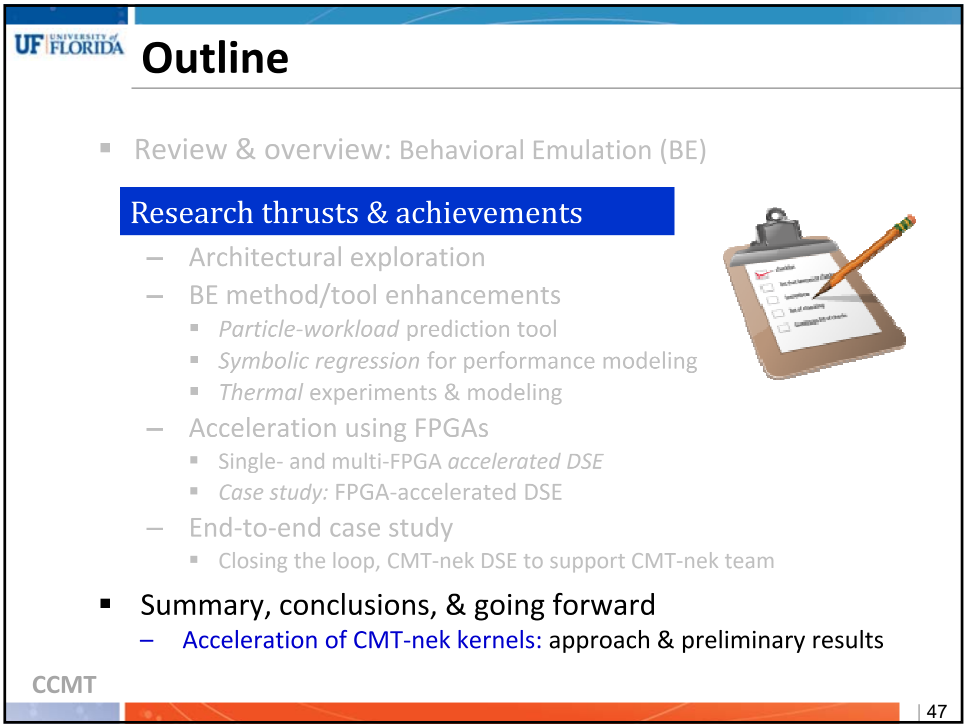

UF UNIVERSITY of FLORIDA

Outline

- Review & overview: Behavioral Emulation (BE)

Research thrusts & achievements

  - Architectural exploration
  - BE method/tool enhancements
    - *Particle-workload* prediction tool
    - *Symbolic regression* for performance modeling
    - *Thermal* experiments & modeling
  - Acceleration using FPGAs
    - Single- and multi-FPGA *accelerated DSE*
    - *Case study:* FPGA-accelerated DSE
  - End-to-end case study
    - Closing the loop, CMT-nek DSE to support CMT-nek team
- Summary, conclusions, & going forward
  - Acceleration of CMT-nek kernels: approach & preliminary results

CCMT

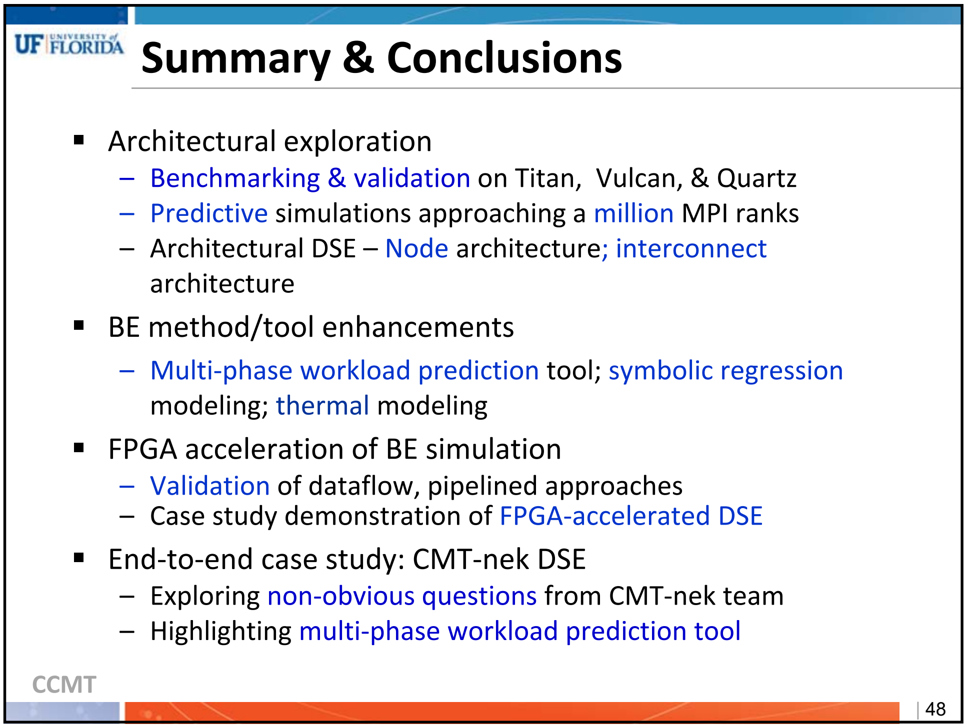

UF UNIVERSITY of FLORIDA

Summary & Conclusions

- Architectural exploration
  - Benchmarking & validation on Titan, Vulcan, & Quartz
  - Predictive simulations approaching a million MPI ranks
  - Architectural DSE – Node architecture; interconnect architecture
- BE method/tool enhancements
  - Multi-phase workload prediction tool; symbolic regression modeling; thermal modeling
- FPGA acceleration of BE simulation
  - Validation of dataflow, pipelined approaches
  - Case study demonstration of FPGA-accelerated DSE
- End-to-end case study: CMT-nek DSE
  - Exploring non-obvious questions from CMT-nek team
  - Highlighting multi-phase workload prediction tool

CCMT

FPGA Acceleration Trends

- FPGAs growing trend in data centers
  - Microsoft Catapult, Amazon F1, Intel Broadwell+FPGA
  - National lab interest
- Recently performed case study on Intel Broadwell+Arria 10 (BDW+A10) for 2D convolution
  - 96x less energy than Broadwell
  - 15.7x less energy than P6000 GPU
- Future research
  - FPGA acceleration of CMT-nek functions
    - Intel Xeon+FPGA system
  - Faculty part of NSF CHREC/SHREC
    - NSF centers for high-performance and low-energy FPGA research

Energy Comparison of BDW+A10 vs. Broadwell

		Image Size				
Precision	*Kernel Size*	*256x256*	*512x512*	*1024x1024*	*2048x2048*	*Avg*
16-bit Fixed	3x3	35×	47×	67×	68×	54×
	5x5	68×	116×	171×	179×	133×
	7x7	58×	135×	153×	164×	128×
	9x9	44×	98×	135×	148×	106×
	Avg	51×	99×	131×	140×	
32-bit Float	3x3	33×	46×	65×	66×	53×
	5x5	64×	111×	162×	169×	127×
	7x7	55×	113×	131×	139×	109×
	9x9	32×	58×	71×	75×	59×
	Avg	46×	82×	108×	112×	

Energy Comparison of BDW+A10 vs. P6000 GPU

		Image Size				
Precision	*Kernel Size*	*256x256*	*512x512*	*1024x1024*	*2048x2048*	*Avg*
16-bit Fixed	3x3	15.5×	19.8×	19.4×	19.0×	18.4×
	5x5	12.9×	18.6×	19.5×	19.8×	17.7×
	7x7	10.5×	18.2×	19.7×	20.5×	17.2×
	9x9	8.1×	15.3×	16.9×	18.6×	14.7×
	Avg	11.8×	18.0×	18.9×	19.5×	
32-bit Float	3x3	14.6×	19.4×	19.0×	18.6×	17.9×
	5x5	12.2×	17.8×	18.6×	18.7×	16.8×
	7x7	9.9×	15.2×	16.9×	17.3×	14.8×
	9x9	5.8×	9.0×	8.9×	9.4×	8.3×
	Avg	10.6×	15.4×	15.8×	16.0×	

CCMT

Do you have any questions?

CCMT

LANL Internship Overview

Frederick Ouellet

Mentor: Cristoph Junghans
Group: CCS-7
ISTI Co-Design Summer School 2018

LA-UR-18-27165

Project Overview

Goal: Optimize performance of the FleCSALE hydrocode and its linking to EOSPAC

Details:

- FlecSALE is a nominally ALE code built on the FleCSI framework
- Supports multiple parallel backends (MPI, Legion, …)
- The EOSPAC library provides the interface from FleCSALE to the SESAME database for equation of state calculations and interpolations
- The test problem run for optimization was the 2D triple point problem

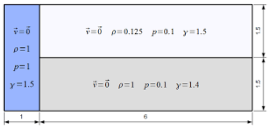

Strategies:

- Optimizing the EOSPAC Interface
- Machine Learning for EoS interpolation
- GPU Porting
- Sparse data optimization (MPI)

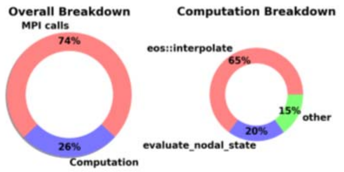

Figure 1: Initial communication and computation breakdown of FleCSALE

CCMT

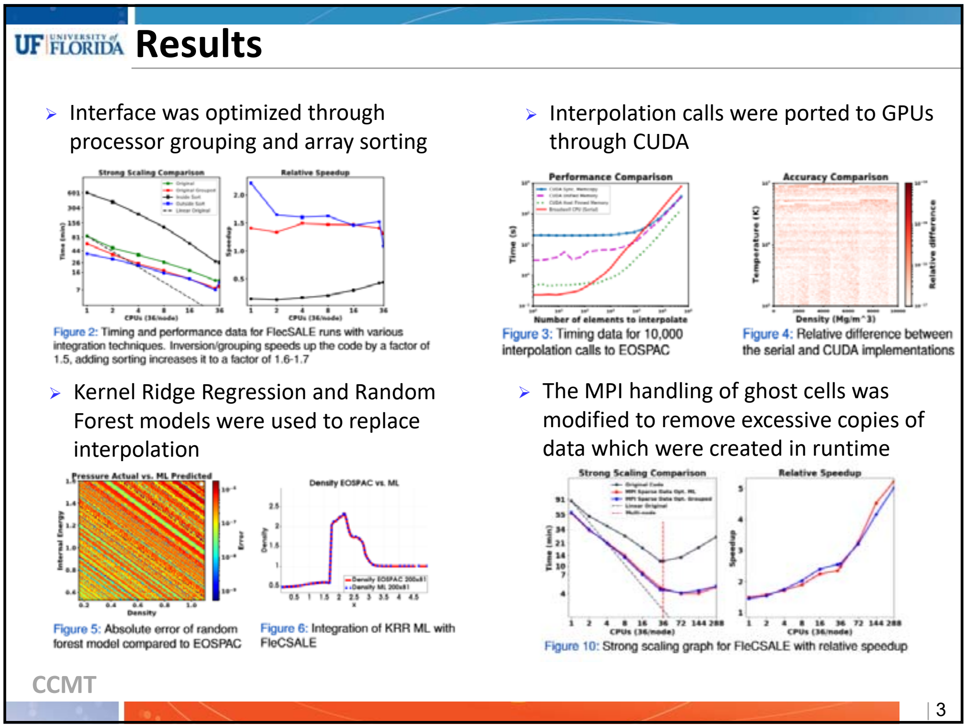

UF UNIVERSITY of FLORIDA

Results

- Interface was optimized through processor grouping and array sorting

Figure 2: Timing and performance data for FlecSALE runs with various integration techniques. Inversion/grouping speeds up the code by a factor of 1.5, adding sorting increases it to a factor of 1.6-1.7

- Interpolation calls were ported to GPUs through CUDA

Figure 3: Timing data for 10,000 interpolation calls to EOSPAC

Figure 4: Relative difference between the serial and CUDA implementations

- Kernel Ridge Regression and Random Forest models were used to replace interpolation

Figure 5: Absolute error of random forest model compared to EOSPAC

Figure 6: Integration of KRR ML with FleCSALE

- The MPI handling of ghost cells was modified to remove excessive copies of data which were created in runtime

Figure 10: Strong scaling graph for FleCSALE with relative speedup

CCMT

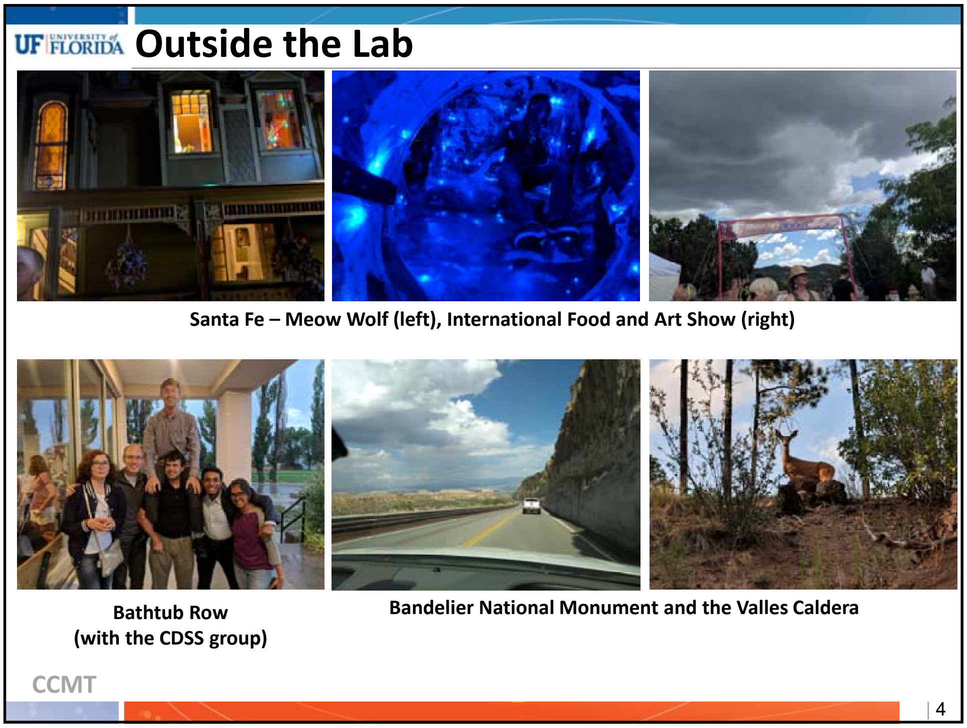

UF UNIVERSITY of FLORIDA

Outside the Lab

Santa Fe – Meow Wolf (left), International Food and Art Show (right)

Bathtub Row
(with the CDSS group)

Bandelier National Monument and the Valles Caldera

CCMT

CCMT

Do you have any questions?