CCMT

CENTER FOR COMPRESSIBLE MULTIPHASE TURBULENCE

AST Review

April 12-13, 2018

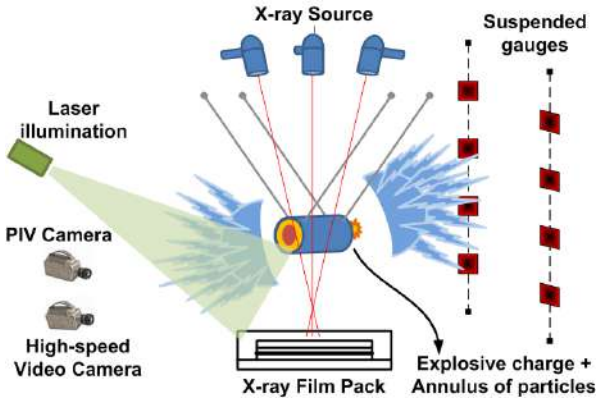

UF UNIVERSITY of FLORIDA

NNSA

UF UNIVERSITY of FLORIDA

Center for Compressible Multi-Phase Turbulence

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

Agenda AST Site Visit April 12-13, 2018

Thursday April 12, 2018

7:45	Van pick up at University Hilton
8:00-9:00	Full Breakfast
	(Review Team and other NNSA personnel will meet in small conference room)
9:00-9:05	Introductions and opening remarks
9:05-10:15	Overview, Y5 Plans, Y4 Accomplishments (Jackson, Balachandar)
10:15-10:30	Coffee break
10:30-11:00	Integration of experiments and simulations (Rollin)
11:00-11:50	CMT-nek (Jason Hackl, David Zwick)
11:50-1:00	Lunch (RT will meet in small conference room)
1:00-1:50	Exascale (Lam, Stitt)
1:50-2:40	CS (Ranka, Banerjee)
2:45-3:00	Coffee break
3:00-3:50	V&V and UQ (Haftka, Park, Kim)
4:00-5:15	Lightning Round of Students' Posters/ Poster Session (light refreshments served)
5:15-6:30	RT Caucus
6:30-8:00	Dinner (Faculty, Staff and Visitors; transportation will be provided for all visitors to the University Hilton)

Friday April 13, 2018

7:45	Van pickup at University Hilton
8:00-9:00	Full Breakfast
	(Review Team and other NNSA personnel will meet in small conference room)
9:00-10:30	Student Presentations – I
	Giselle Fernandez (15 mins) Chandler Moore (15 mins) Goran Marjanovic (15 mins) Kyle Hughes (15 mins) Heather Zunino (ASU; 30mins)
10:30-10:45	Coffee Break
10:45-11:15	Student Presentations – II
	Keke Zhai (15 mins) Sai Chenna (15 mins)
11:15-12:15	Y6 Plans and Center Response to RT Questions (PI Team)
12:15-4:00	Lunch (RT will meet in small conference room)
	Private RT deliberations (small conference room)
	Discussions between Center Management and AST as appropriate (large conference room)
4:00-4:30	RT Summary for Center Management (large conf. room)
4:30	Review ends

Center for Compressible Multiphase Turbulence

AST Review April 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

Sam Schofield (Chair)	LLNL	schofield5@llnl.gov
Brian Carnes (V&V/UQ)	SNL	bcarnes@sandia.gov
Robert Clay (CS)	SNL	rlclay@sandia.gov
Gabe Rockefeller (CS)	LANL	gaber@lanl.gov
Kambiz Salari (Physics)	LLNL	salari1@llnl.gov
Erik Vold (Physics)	LANL	elv@lanl.gov
Jeremiah Wilke (CS)	SNL	jjwilke@sandia.gov

Others

Tina Macaluso	Leidos/NESD	antoinette.macaluso@leidos.com
Bob Voigt	Leidos/NESD	rvoigt@krellinst.org
John Feddema	SNL	jtfedde@sandia.gov
Fernando Grinstein	LANL	fgrinstein@lanl.gov
Ana Kupresanin	LLNL	kupresanin1@llnl.gov
Alan Kuhl	LLNL	kuhl2@llnl.gov
Brian Taylor	Eglin AFB	brian.taylor.56@us.af.mil
Greg Weirs	Sandia	vgweirs@sandia.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

Center for Compressible Multiphase Turbulence

Students

Ryan Blanchard	University of Florida	rlb1116@ufl.edu
Sai Chenna	University of Florida	chenna@hcs.ufl.edu
Paul Crittenden	University of Florida	pcritte@ufl.edu
Brad Durant	University of Florida	neoncrash@ufl.edu
Giselle Fernandez	University of Florida	gisellefernandez@ufl.edu
Mohamed Gadou	University of Florida	mgadou@ufl.edu
Joshua Garno	University of Florida	jgarno@ufl.edu
Trokon Johnson	University of Florida	johnson@hcs.ufl.edu
Tanner Jones	University of Florida	tannerwaynejones@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
Goran Marjanovic	University of Florida	gmarjanovic@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
Brandon Osborne	University of Florida	bosborne@ufl.edu
Frederick Ouellet	University of Florida	f.ouellet@ufl.edu
Carlo Pascoe	University of Florida	poppyc@ufl.edu
Raj Rajagoplan	University of Florida	saravana8394@ufl.edu
Ben Reynolds	University of Florida	beynoldsren@gmail.com
Chad Saunders	University of Florida	csaunders17@ufl.edu
Shirly Spath	University of Florida	spathsy@ufl.edu
Prashanth Sridharan	University of Florida	shan1130@ufl.edu
Cameron Stewart	University of Florida	csstewart10@ufl.edu
Keke Zhai	University of Florida	zhaikeke@ufl.edu
Yiming Zhang	University of Florida	yiming90521@gmail.com
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 & Integration
Y4 Accomplishments
Y5 Plans

T.L. Jackson

S. Balachandar

UF UNIVERSITY of FLORIDA

AST Meeting Agenda

Thursday

- Overview, Integration, Y5 Plans, Y4 Accomplishments (Jackson, Balachandar)
- Integrated Simulations: Transitioning from RocFlu to CMT-nek (Rollin)
- CMT-nek (Hackl, Zwick)
- Lunch
- BE Results at Scale (Lam, Stitt)
- CMT-nek: CS Updates (Ranka)
- V&V and UQ (Park)
- Student Lightning Round
- Poster Session
- RT Caucus
- Dinner 6:30-8:00 (University Hilton; Dogwood Room)

Friday

- Student Presentations (7)
- Y6 Plans and Response to RT Questions (PI Team)
- Lunch
- RT Deliberations/Summary

CCMT 2

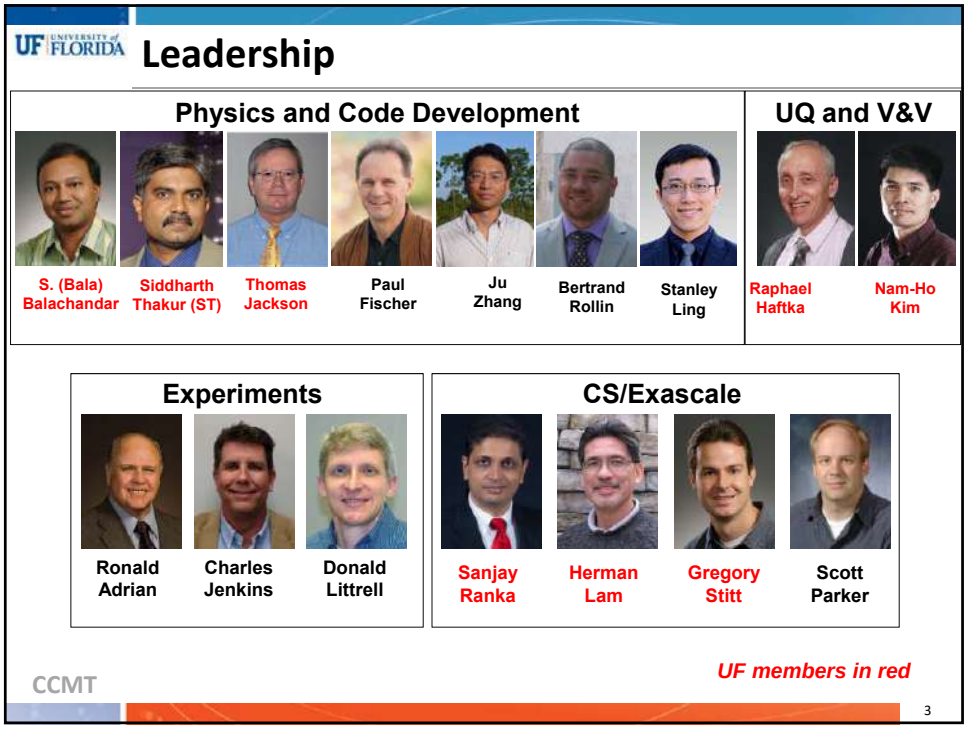
UF UNIVERSITY of FLORIDA
Leadership
Physics and Code Development
S. (Bala) Balachandar
Siddharth Thakur (ST)
Thomas Jackson
Paul Fischer
Ju Zhang
Bertrand Rollin
Stanley Ling
UQ and V&V
Raphael Haftka
Nam-Ho Kim
Experiments
Ronald Adrian
Charles Jenkins
Donald Littrell
CS/Exascale
Sanjay Ranka
Herman Lam
Gregory Stitt
Scott Parker
UF members in red
CCMT
3


UF UNIVERSITY of FLORIDA
Research Staff & Senior PhD Students
Tania Banerjee
Angela Diggs (Eglin AFB)
Kei Fujisawa
Jason Hackl
Nguyen T. Nguyen
Chanyoung Park
Carlo Pascoe
CCMT
4

Current Students (Undergraduate & Graduate)

CCMT

5

Internship Program – Completed (16)

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

CCMT

6

Internship Program – Completed (16)

Goran Marjanovic	Sandia	Aug-Nov, 2016	Drs. Paul Crozier & Stefan Domino
Georges Akiki	LANL	May-Aug, 2016	Dr. Marianne Francois
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

CCMT 7

Internship Program – Not Completed

- Summer 2018
  - Prashanth Sridharan (MAE, Physics); LANL
  - Fred Ouellet, PhD (MAE, Physics and UQ); LANL
- Brad Durant, PhD (MAE, Physics and UQ)
- Joshua Garno, PhD (MAE, Physics and UQ)
- Chandler Moore, PhD (MAE, Physics) (NSF Fellowship)

CCMT 8

Graduated Students & Postdocs

- 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
- Bertrand Rollin, Postdoc in thru August 2014
  - Assistant Professor, Embry Riddle, Daytona Beach FL
- Mrugesh Shringarpure, Postdoc in thru January 2016
  - Research Engineer, ExxonMobil, Houston TX
- Subbu Annamalai, PhD (2015), Dr. S. Balachandar, MAE; Postdoc in thru March 2017
  - Senior Systems Engineer, Optym, Gainesville FL
- Parth Shah, MS (2016), Drs. H. Lam and G. Stitt
- Georges Akiki, PhD (2016), Dr. S. Balachandar, MAE; Postdoc thru March 2017
  - Postdoctoral Associate, LANL

CCMT 9

Graduated Students & Postdocs

- 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
- Yiming Zhang, PhD (2018), Drs. Haftka and Kim,MAE
  - GE Global Research, Niskayuna, NY

CCMT 10

PhD Students Expected Graduation Dates

- 2018
  - Paul Crittenden
  - Giselle Fernandez
  - Mohamed Gadou
  - Kyle Hughes
  - Goran Marjanovic
  - Yash Mehta
  - Carlo Pascoe
  - Prashanth Sridharan
  - Heather Zunino
- 2019
  - Rahul Koneru
  - Samaun Nili
  - Fred Ouellet
  - Keke Zhai
  - David Zwick
- 2020
  - Brad Durant
  - Josh Garno
- 2021
  - Sai Chenna
  - Trokon Johnson
  - Adeesha Malavir
  - Chandler Moore
  - Aravind Neelakantan
  - Raj Rajagoplan

CCMT 11

Additional Information

- Additional Graduate Program Announcements
  - David Zwick – NSF Fellowship Graduate Program (Aug 2016)
  - Georges Akiki - MAE Best Dissertation Award (TSFD; May 2017)
  - Chandler Moore – NSF Fellowship Graduate Program (Aug 2017)
- Other metrics (Y1 – Y4)
  - Publications: 134
  - Presentations: 94
- Deep Dive Workshops
  - Exascale & CS Issues, Feb 3-4, 2015, University of Florida
  - Multiphase Physics, Oct 13-14, 2016, Tampa FL
  - CMT-nek/nek5000, April 17-18, 2018, Tampa FL
  - Multiphase Deep Dive 2; Stanford lead with Florida co-lead; Fall 2018
- Center Webpage
  - http://www.eng.ufl.edu/ccmt/

CCMT 12

Additional Information

- CMT-nek/nek5000
  - April 17-18, 2018, Tampa FL
  - 24 talks
    - 4 from Florida – expanding nek5000
  - 33 participants
    - DOE laboratories represented
    - 3 Countries (US, Canada, Sweden)

CCMT

13

Additional Information

- 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 Work Camp, 2017

CCMT

14

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 – new 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)

CCMT 15

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

CCMT 16

4th Annual CCMT Picnic (March, 2018)

CCMT

17

Outline

- Background
- Scope of the center
- Y4 accomplishments
- Y5 plans
- RocFlu to CMT-nek transition
- Integration and timeline

CCMT

18

Demonstration Problem

CCMT

19

Prediction Metrics

PM-1: Blast Wave Location

PM-2: Particle Front Location

PM-3: Number of Instability Waves

PM-4: Amplitude of Instability Waves

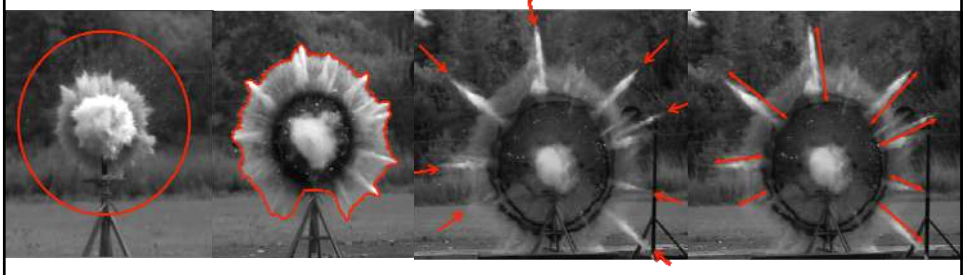

CCMT

20

Sequence of Events

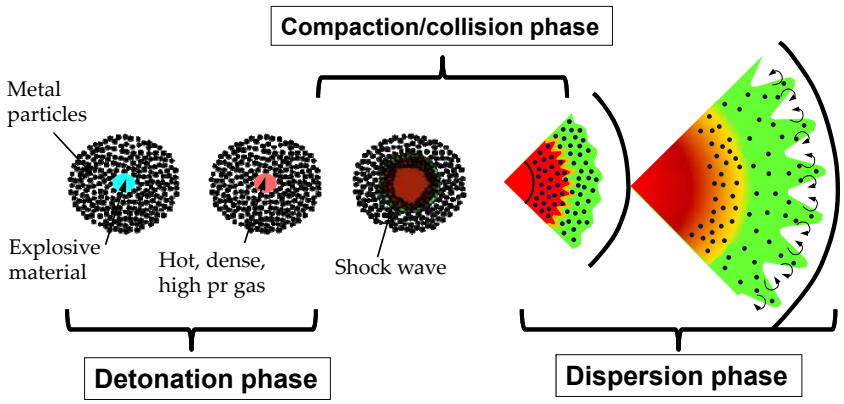

CCMT

21

Physical Models – Sources of Error

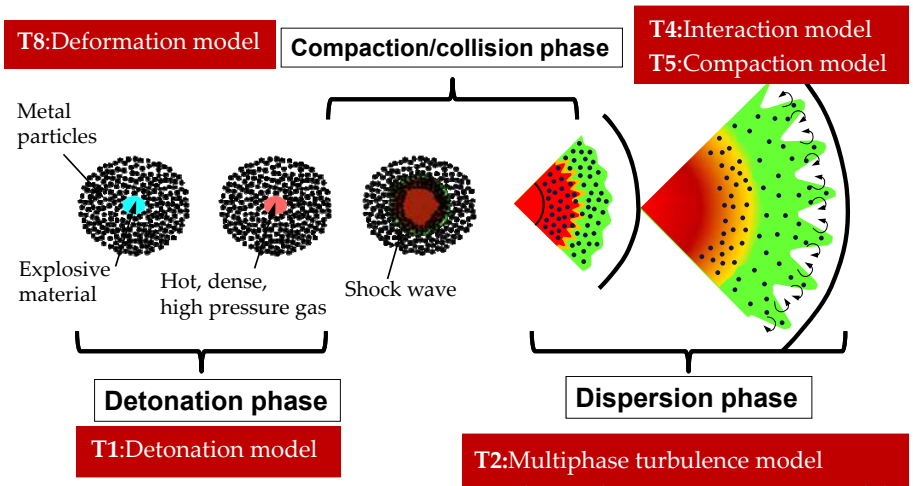

CCMT

22

Multiscale Integration Strategy

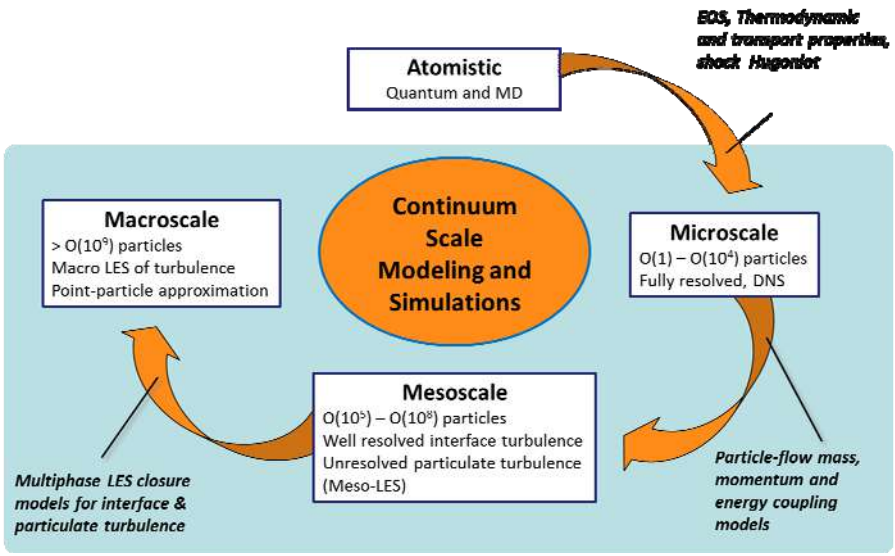

CCMT 23

Updated Scope

- Our focus will be on
  - Rayleigh-Taylor (RT) and Richtmeyer-Meshkov (RM) instabilities
  - Gas-particle coupling
  - Mixing at the rapidly expanding material front
  - Self-assembly of explosive-driven particles
  - Reduced emphasis on initial compaction
- Will avoid the following complications
  - Free-shear and wall turbulence (stay away from boundaries)
  - Detonation physics (use simple, well-studied explosives)
  - Fragmentation or atomization physics (avoid casing, liquids)
  - Reactive physics (use non-reactive metal particles)

CCMT 24

Sources of Errors & Uncertainties

- T1: Detonation modeling
- T2: Multiphase turbulence modeling
- T3: Thermodynamics & transport properties
- T4: Particle-particle interaction modeling
- T5: Compaction modeling
- T6: Force coupling modeling
- T7: Thermal coupling modeling
- T8: Particle deformation and other complex physics

- T9: Discretization and numerical approximation errors
- T10: Experimental and measurement errors & uncertainties

Key Focus

Advance state-of-the-art

- Multiphase turbulence

- Force coupling model

CCMT 25

Uncertainty Budget – Overall Plan

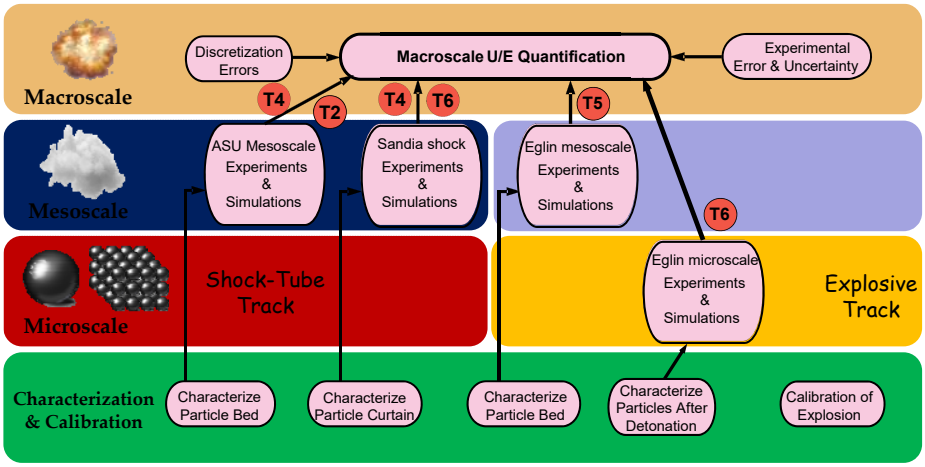

- T2 – Turbulence modeling
- T4 – Particle interaction modeling
- T5 – Compaction modeling
- T6 – Force coupling modeling

CCMT 26

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

CCMT 27

Uncertainty Reduction Workflow

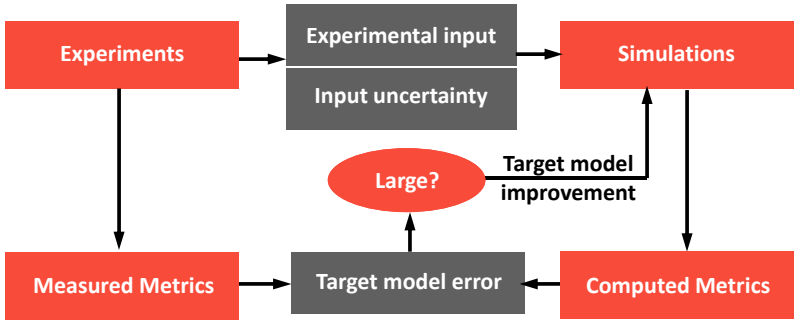

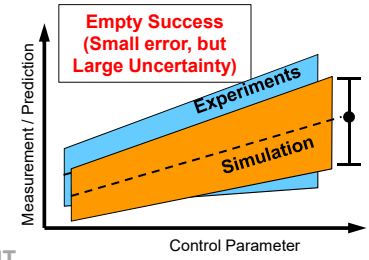

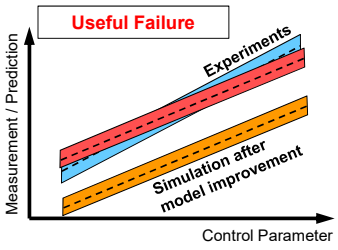

CCMT 28

Uncertainty Budget Workflow

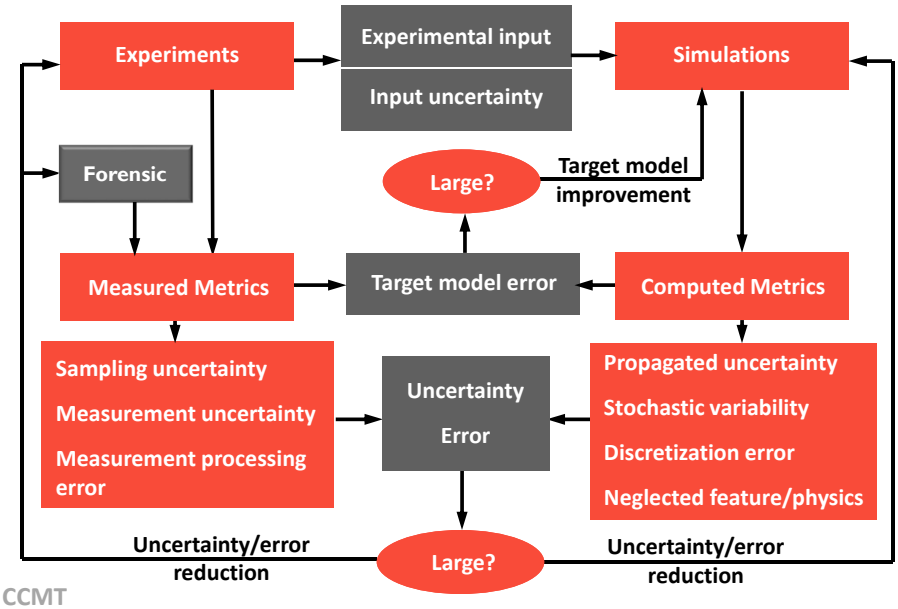

CCMT

29

UB Workflow - Experiment Worksheet

- Experimental input
  - Shock properties, particle properties, curtain properties, …
- Input uncertainty
  - Quantified uncertainties in all the above
- Prediction metrics
  - PM1: Shock position, PM2: Upstream and downstream curtain
- Uncertainty & error quantification (UQ)
  - Error in PM1 and PM2 obtained from Schlieren
  - Error in X-ray image particle volume fraction
- Uncertainty & error reduction (UR)
  - Perform new experiments without spanwise gap
  - Improved measurement to reduce volume fraction uncertainty

CCMT

30

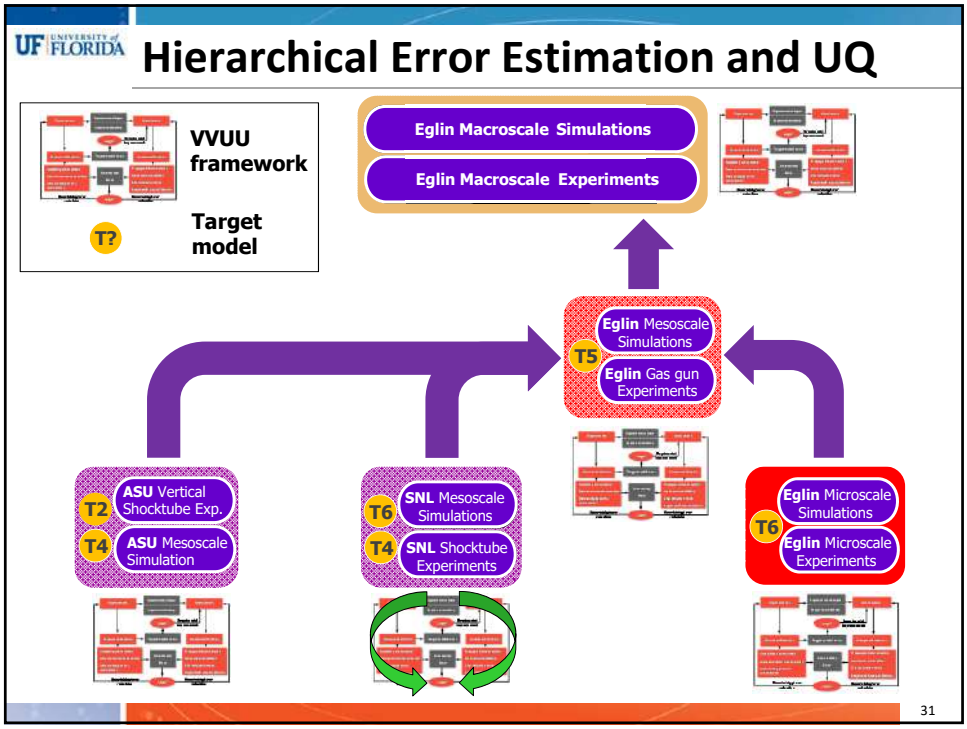
UF UNIVERSITY of FLORIDA
Hierarchical Error Estimation and UQ
VVUU framework
T?
Target model
Eglin Macroscale Simulations
Eglin Macroscale Experiments
T5
Eglin Mesoscale Simulations
Eglin Gas gun Experiments
T2
ASU Vertical Shocktube Exp.
T4
ASU Mesoscale Simulation
T6
SNL Mesoscale Simulations
T4
SNL Shocktube Experiments
T6
Eglin Microscale Simulations
Eglin Microscale Experiments
31

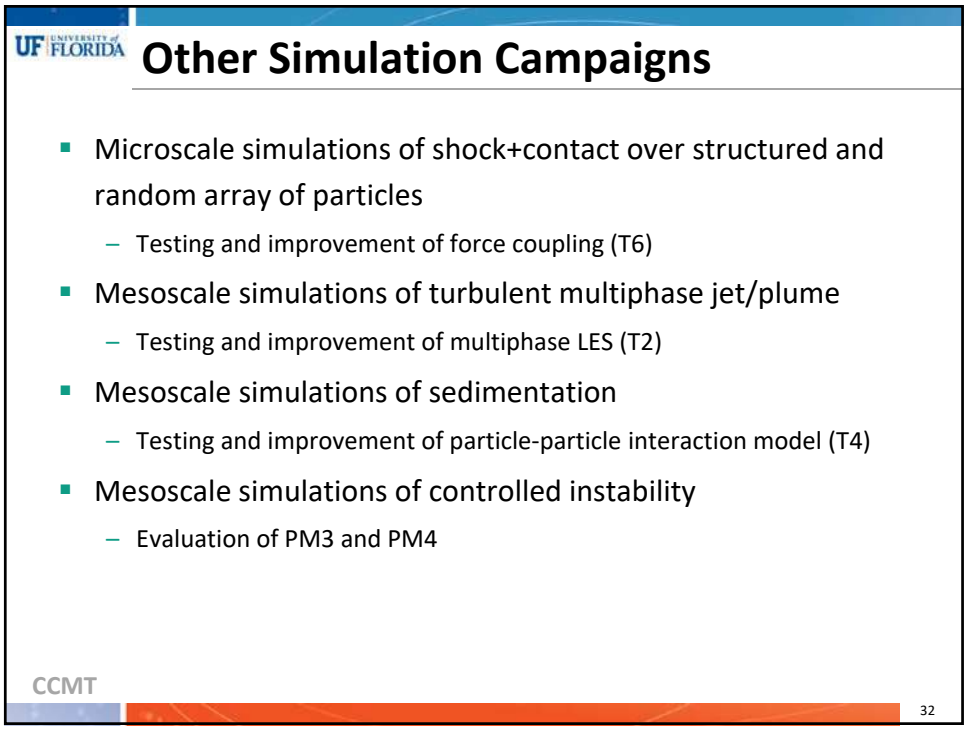
UF UNIVERSITY of FLORIDA
Other Simulation Campaigns
Microscale simulations of shock+contact over structured and random array of particles
Testing and improvement of force coupling (T6)
Mesoscale simulations of turbulent multiphase jet/plume
Testing and improvement of multiphase LES (T2)
Mesoscale simulations of sedimentation
Testing and improvement of particle-particle interaction model (T4)
Mesoscale simulations of controlled instability
Evaluation of PM3 and PM4
CCMT
32

Key Accomplishments

- Tight integration
- Empower students and staff

- PIEP model + Machine Learning
- CMT-nek a versatile multiphase flow code (scalable to >O(10^6) core)
- Culture of UQ integration
- BE framework and FPGA
- Dynamic load balancing

CCMT

33

Y4 Highlights

1. Macroscale – Hero Run
2. Blastpad & other validation experiments
3. CMT-nek development and transition
4. Mesoscale – CMT-nek simulation of expansion fan
5. Microscale – Shock + Contact
6. UQ workflow
7. Design space exploration with Behavioral Emulation
8. Dynamic load balancing of Euler-Lagrange

CCMT

34

1: Demonstration Problem (Macroscale)

Goal

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

Year 4

- Used existing code to perform largescale simulations of the demonstration problem
- Qualitative comparison against experimental data of Frost (PM1 & PM2)
- Integrate additional capabilities for the hero runs: real gas EOS, reactive burn

Presentation

- Bertrand Rollin

Hero-3 demonstration simulation (Rocflu)

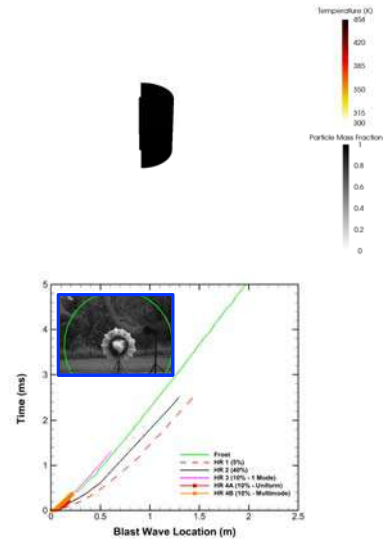

CCMT

35

1: Demonstration Problem (Macroscale)

Goal

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

Year 4

- Used existing code to perform largescale simulations of the demonstration problem
- Qualitative comparison against experimental data of Frost (PM1 & PM2)
- Integrate additional capabilities for the hero runs: real gas EOS, reactive burn

Presentation

- Bertrand Rollin

CMT-Hero-1 demonstration simulation

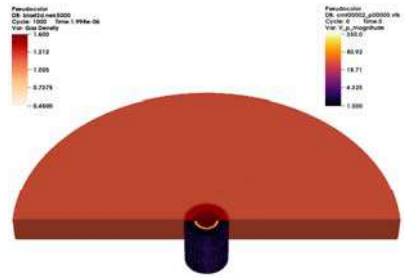

- Features:
  - 32.4 M Degrees of freedom
  - 1 M computational particles
  - 0.3 x 0.3 x 0.0015 m
  - t_{max} = 0.3ms
  - 32768 MPI ranks

CCMT

36

2: Blastpad Experiments

Goals

- Obtain validation-quality experimental measurements of the demonstration problem
- Validation-quality experiments at micro and mesoscales
- Perform shock-tube track micro- and mesoscale experiments

Year 4

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

Presentation

- Bertrand Rollin, Kyle Hughes

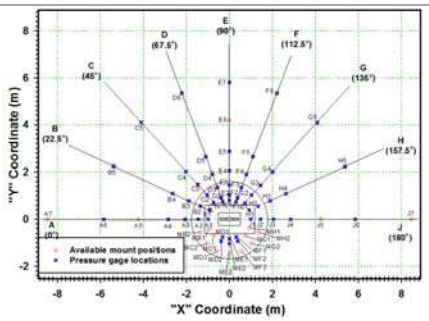

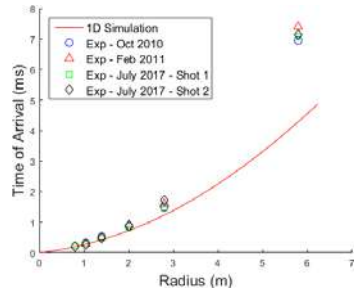

CCMT

37

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.

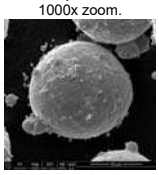

SEM of several steel particles at 100x zoom.

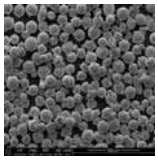

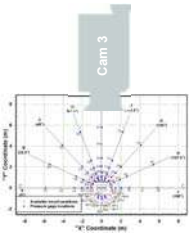

CCMT

38

3: CMT-nek Development

Goals

- Co-design an exascale code (CMT-nek) for compressible multiphase turbulence
- Perform micro, meso and demonstration-scale simulations
- Develop & incorporate energy and thermal efficient exascale algorithms

Year 4

- Developed and released microscale version of CMT-nek for microscale simulations
- Developed and released mesoscale version of CMT-nek for mesoscale simulations
- Shock capturing with EVM
- CMT-nek in nek5000 repository

Presentation

- Jason Hackl and David Zwick
- Goran Marjanovic

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

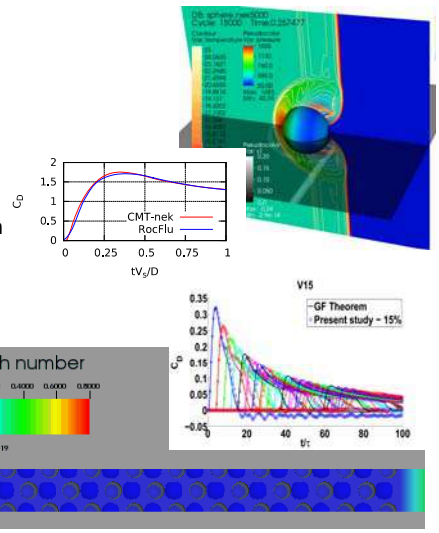

CCMT

39

Multiphase Capabilities

			nek5000	CMT-nek
Single-phase	Fluid-Fluid	Incompressible	✓	
		Low-Mach-number	✓	✓
		Shock waves		✓
Increased multiphase coupling	Tracer Particles		✓	✓
	Particle-Fluid	Dilute	✓	✓
		Dense	✓	✓
	Particle-Particle		✓	✓

✓ Existed ✓ New

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

CCMT

40

Algorithmic Scaling

Ideal Setup:

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

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

Strong Scaling (s):

- 15 million grid points
- 8.4 million particles

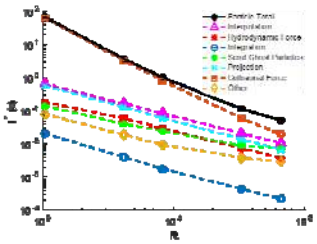

Weak Scaling (w):

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

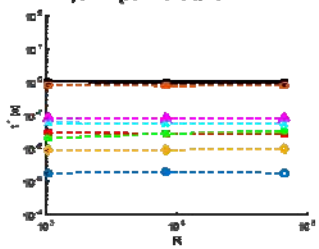

- **Presentation**
  - Keke Zhai

CCMT

41

4: Mesoscale Simulations & Experiments

Goal

- Perform a hierarchy of mesoscale simulations to allow rigorous validation, uncertainty quantification and propagation to the demonstration problem

Year 4

- 4-way coupled simulation with CMT-nek
- Mesoscale simulations of expansion fan over a bed of particles
- 4-way coupled simulations with and without PIEP
- Bundled simulations for UQ

Presentation

- David Zwick

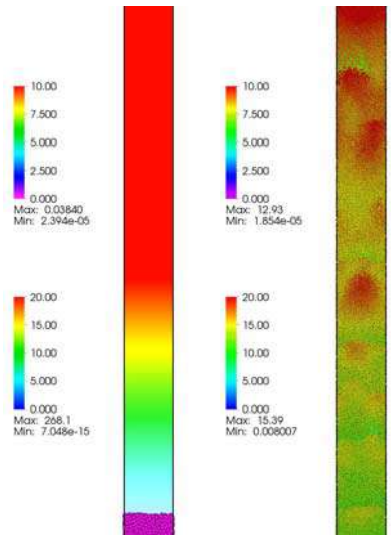

Case 1 Case 2

CCMT

42

4: Mesoscale Simulations & Experiments

Goals

- Obtain validation-quality experimental measurements of the demonstration problem
- Validation-quality experiments at micro and mesoscales
- Perform shock-tube track micro- and mesoscale experiments

Year 4

- Expansion fan experiments at ASU
- Detailed instrumentation for validation
- Simulation informed experiments
- Integrated UQ

Presentation

- Heather Zunino (ASU)

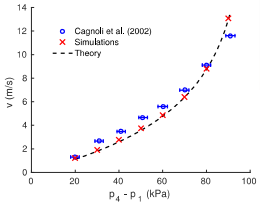

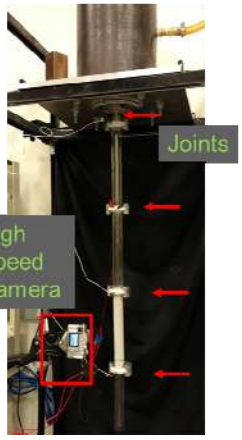

CCMT

43

5: Microscale Simulations

Goals

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

Year 4

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

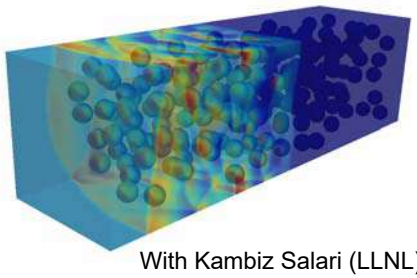

With Kambiz Salari (LLNL)

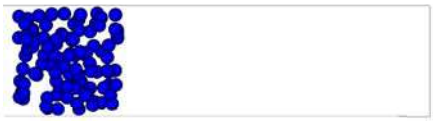

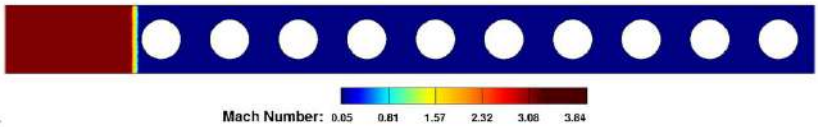

CCMT

44

5: Pairwise Interaction Extended Point-Particle Model

$\phi = 44\%, Re = 20, N = 459$

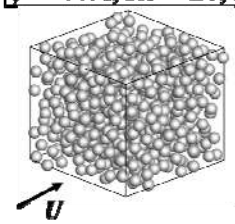

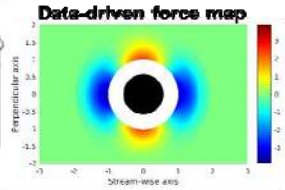

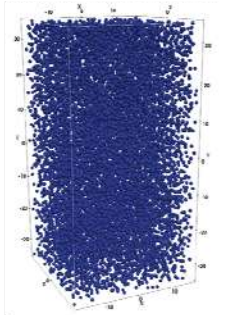

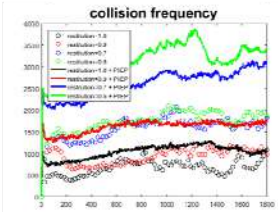

R^2 Values:		PIEP-Physics			PIEP-Hybrid		
ϕ	Re	Drag	Lift	Torque	Drag	Lift	Torque
0.1	40	0.70	0.68	0.75	0.72	0.75	0.79
0.1	70	0.65	0.68	0.65	0.67	0.72	0.72
0.1	173	0.35	0.59	0.45	0.44	0.60	0.58
0.2	10	0.38	0.34	0.48	0.66	0.74	0.71
0.2	80	0.50	0.48	0.72	0.62	0.63	0.77
0.45	21	0.01	0.09	0.47	0.55	0.59	0.76
0.45	116	0.21	0.19	0.51	0.63	0.57	0.65

- **Presentation**: Chandler Moore

CCMT

45

6: UQ Workflow

Goals

- Develop UB as the backbone of the Center
- Unified application of UB for both physics and exascale emulation

Year 4

- Identify main uncertainty sources and quantify their contributions to the model uncertainty of the shock tube simulation
- Reduce uncertainty to focus on model error
- UQ and propagation in the context of exascale emulation
- JWL mixture-EOS surrogate for efficient computation

Presentation

- Nam-Ho Kim, Giselle Fernandez, Kyle Hughes

Sandia shock tube UQ Workflow

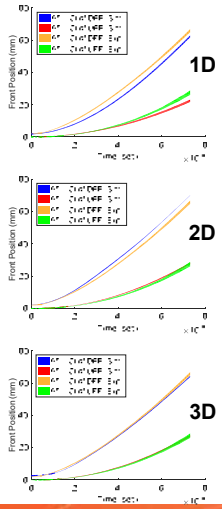

Pseudo-turbulence

CCMT

46

UQ of Eglin Microscale

U in Density with CI 95%
U in Diameter with CI 95%
U in Explosive Density with CI 95%

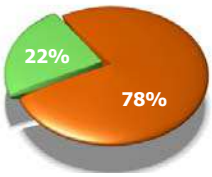

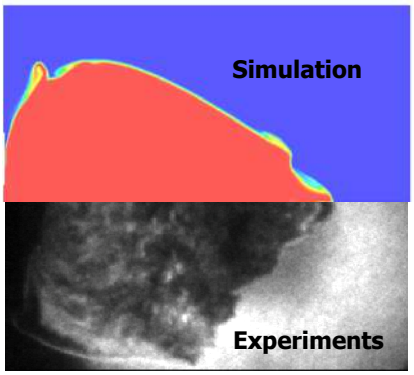

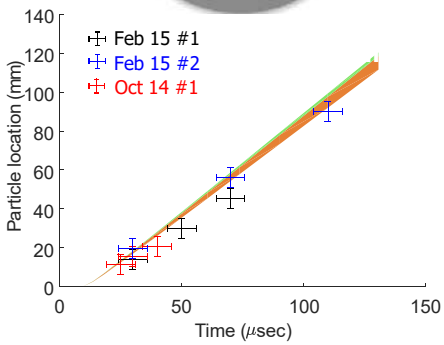

47

7: Exascale Emulation

Goal

- Develop behavioral emulation (BE) methods and tools to support co-design for algorithmic design-space exploration and optimization of key CMT-bone kernels & applications on future Exascale architectures

Year 4

- Enhanced BE methods with network models, interpolation schemes, and benchmarking for CMT-bone AppBEOs
- Performed large-scale experiments on DOE platforms with BE-SST simulator
- Started design space exploration
- Improved throughput and scalability for FPGAs

Presentation

- Herman Lam, Greg Stitt, Sai Chenna

Nek-5000 on KNL (64 cores)

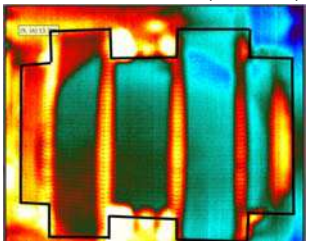

Nek-5000 on KNL (1 core)

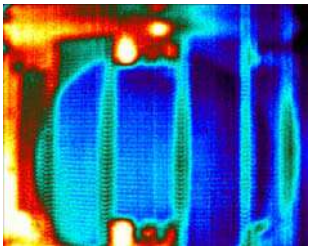

CCMT

48

Design Space Exploration: Results

- **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)

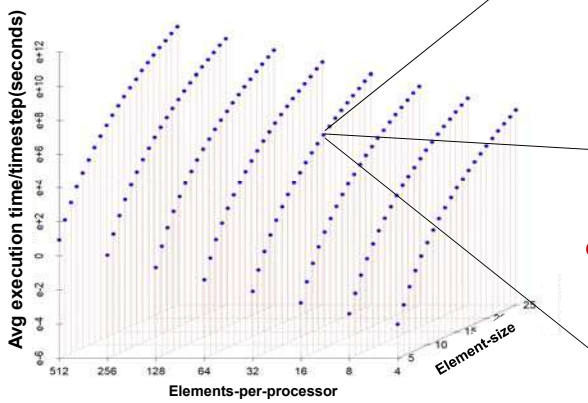

BE-FPGA Simulation

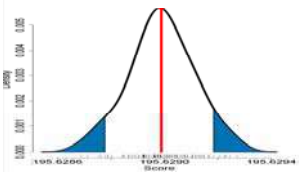

BE-SST Simulation (32k cores)

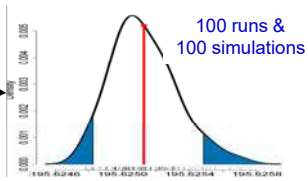

CMT-nek Particle-Solver Execution (4k cores)

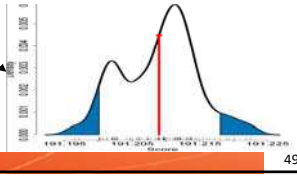

CCMT *ran on HiPerGator - 64 cores **100 Monte-Carlo simulations/candidate

49

8: Dynamic Load Balancing

Goal

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

Year 4

- Carried out extensive investigation of performance and energy issues for CMT-bone
- Hybrid CPU-GPU implementation of CMT-bone and optimization
- Thermal aware optimization
- Dynamic load balancing with particles

Presentation

- Sanjay Ranka, 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

50

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

51

CMT-Hero Timeline

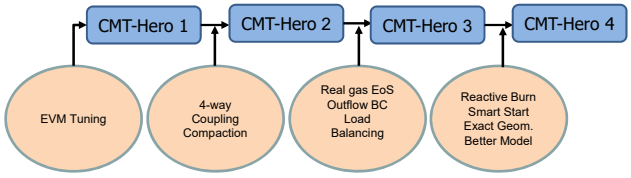

		Y4-Q1	Y4-Q2	Y4-Q3	Y4-Q4	Y5-Q1	Y5-Q2	Y5-Q3	Y5-Q4
S4a	**CMT-Hero 1**								
	EVM tuning								
S4b	**CMT-Hero 2**								
	Compaction								
	4-way coupling								
S4c	**CMT-Hero 3**								
	Real gas EOS								
	Outflow BC								
	Load Balancing								
S4d	**CMT-Hero 4**								
	Reactive burn								
	Smart start								
	Improved models								

CCMT

52

UF UNIVERSITY of FLORIDA

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		

CCMT

53

UF UNIVERSITY of FLORIDA

Transition to CMT-nek – Simulation Plan

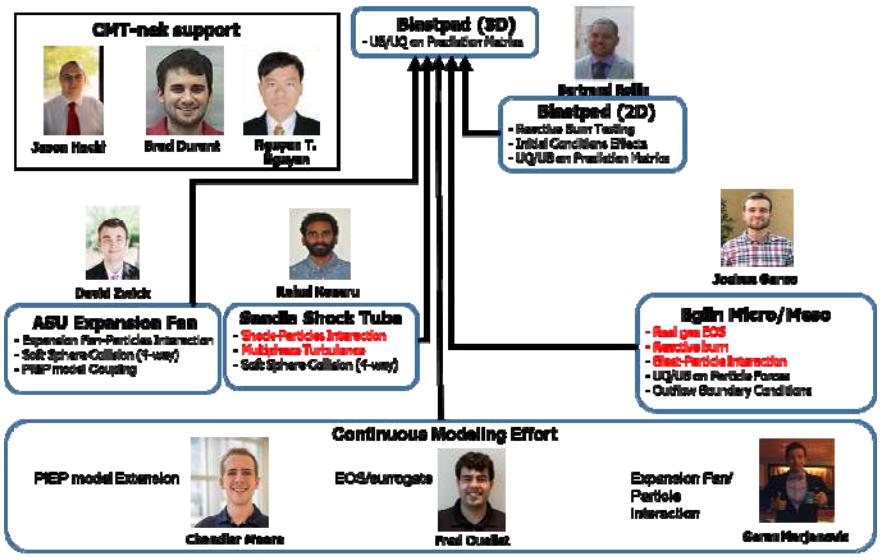

CCMT

54

CMT-nek Micro-Meso-Hero Timeline

		Y4-Q1	Y4-Q2	Y4-Q3	Y4-Q4	Y5-Q1	Y5-Q2	Y5-Q3	Y5-Q4	
S1	**Micro Sims**									➢ Goran Marjanovic
	4-way coupling									
	Load Balancing									
	Compaction									
S2	**ASU Meso**									➢ David Zwick
	Outflow BC									
	Improved models									
S3	**Eglin Meso**									➢ Josh Garno
	EVM tuning									
	Smart start									
	Real gas EOS									
	Reactive Burn									
S4	**CMT-Hero**									➢ Rahul Koneru ➢ Fred Ouellet

CCMT

55

Major Plans for Y5

- Complete CMT-Hero2, CMT-Hero3, and CMT-Hero4
- Complete a UB cycle for ASU expansion fan, Eglin mesoscale
- Propagate uncertainties and perform the UB cycle for the Eglin blastpad experiment
- Complete ASU simulation at an unprecedented resolution and detailed comparison with experimental measurements
- Finalize load balancing and release it for wide external use
- Complete the implementation of CMT-nek on CPU-GPU
- Predict and later test the performance of CMT-nek on "Summit"
- End-to-end design space exploration of CMT-nek using multiple FPGA

CCMT

56

Risks & Mitigation

- Comprehensive risk assessment and backup plans from each team
- We have successfully crossed some fundamental risks we started with: Will DGSEM work for compressible flows with shocks? Will Eglin carry out macroscale experiments?
- Optimal tuning of EVM for extreme blast conditions
  - Stay with first order wave-speed viscosity, etc.
- Particle compaction is too hard at explosive condition
  - Solve very early time using 2-fluid model
- Building even a surrogate of the macroscale problem may be too expensive
  - Limit to one or two high fidelity & several lower fidelity simulations
- Current load balancing strategy may face memory limitation
  - Experiment domain partitioning strategies that loosely maintain spatial locality
- BE-SST simulation of beyond 1M MPI ranks
  - Limit to 1M MPI ranks

CCMT

57

CCMT

Do you have any questions?

CCMT

Integrated Simulations:

Transitioning from Rocflu To CMT-nek

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


Brandon Osborne


Prashanth Sridharan


Paul Crittenden


Chandler Moore

CCMT

2

From Rocflu to CMT-nek

	Rocflu	CMT-nek
Method	Finite volume	Discontinuous Galerkin
Mesh	Mixed cells (tetrahedra, hex)	Spectral elements (hex)
Shock capturing	WENO gradient reconstruction	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, 11:00 am

CCMT

3

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

4

Explosive-Driven Particle-Laden Flows

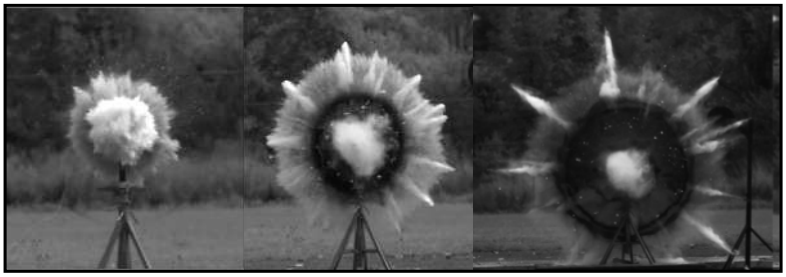

- Bring predictive capabilities to particle-laden flow simulations under extreme conditions

CCMT

5

Sequence of Events and Physical Modeling

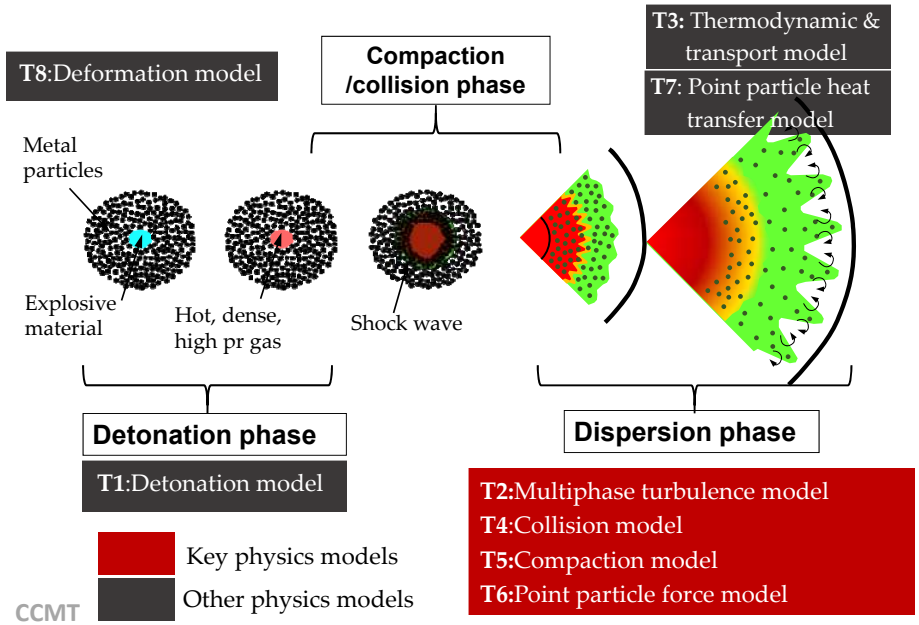

CCMT

6

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		

CCMT

7

CMT-nek Development

- EVM parameter tuning
- Equation of State/Multi-species formulation
- Reactive Burn Initial Conditions
- Boundary Conditions
- 4-way coupling
- Particles Compaction
- Particles Load Balancing
- Intelligent Starting Condition
- Integration of improved coupling models

➢ Jason Hackl and David Zwick, 11:00 am

CCMT

8

CMT-Hero Timeline

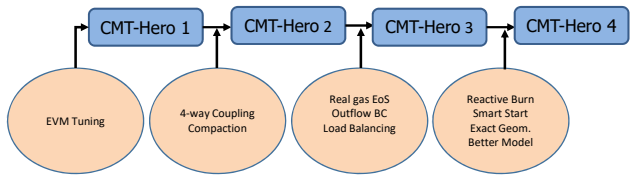

		Y4-Q1	Y4-Q2	Y4-Q3	Y4-Q4	Y5-Q1	Y5-Q2	Y5-Q3	Y5-Q4
S4a	**CMT-Hero 1**								
	EVM tuning								
S4b	**CMT-Hero 2**								
	Compaction								
	4-way coupling								
S4c	**CMT-Hero 3**								
	Real gas EOS								
	Outflow BC								
	Load Balancing								
S4d	**CMT-Hero 4**								
	Reactive burn								
	Smart start								
	Improved models								

CCMT

9

CMT-HERO 1 – Initial Conditions

Mesh Parameters	
Domain size (L_x x L_y x L_z)	0.3m x 0.3m x 0.015m
Number of elements (N_x x N_y x N_z)	1.2 M (636 x 636 x 3)
Order of polynomial (p)	3
Number of Degree of Freedom	32.4 M

Initial Conditions		
Quantity	**Charge**	**Ambient**
Pressure	100 atm	1 atm
Density	23.53658 kg/m^3	1.20432 kg/m^3
Temperature	1500 K	293.15 K

Charge radius	3.8 mm
Particle bed thickness	2 cm
No. of particles	1 M
Particle Volume Fraction	5%

Parameters in Entropy Viscosity Method	
C_{max} (Constant in wave speed based viscosity)	1.0
C_E (Constant in entropy residual based viscosity)	40

CCMT

10

CMT-Hero 1

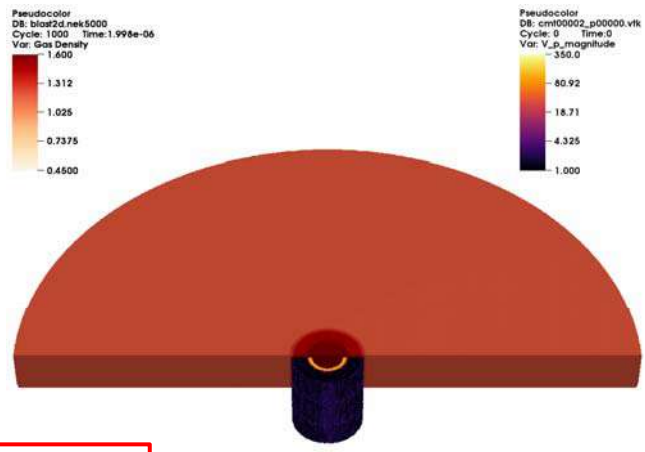

- Features:
  - 32.4 M Degrees of freedom
  - 1 M computational particles
  - 0.3 x 0.3 x 0.0015 m
  - t_{max} = 0.3ms
  - 32768 MPI ranks

➢ Chanyoung Park, 3:00 pm

11

CMT-Hero 1 – Profiles

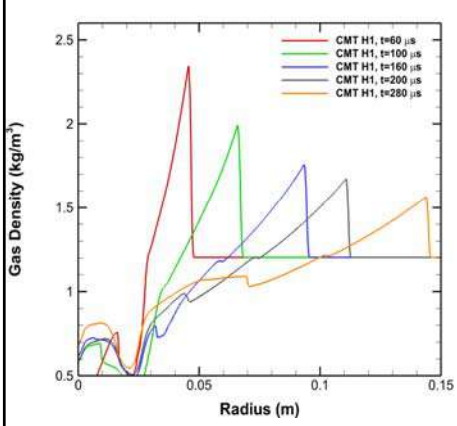

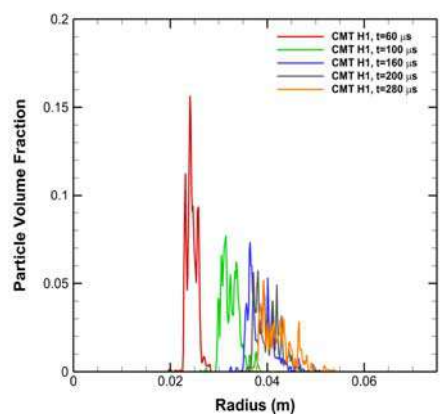

CCMT

12

CMT-Hero 1 – Prediction Metrics

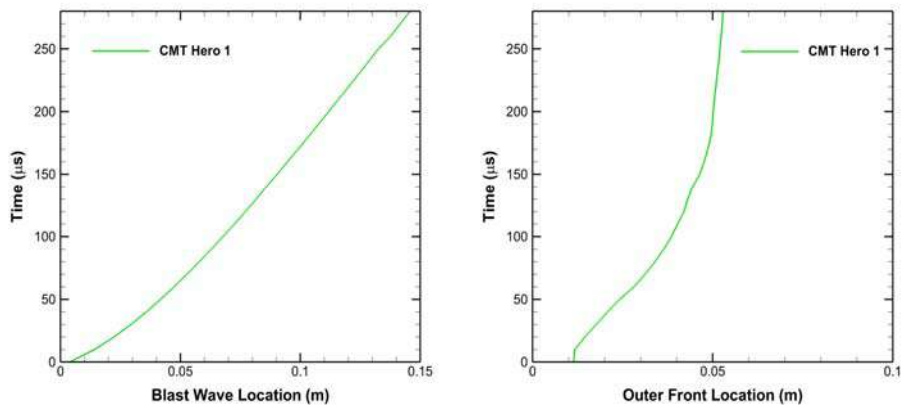

CCMT

13

Hero Run 3

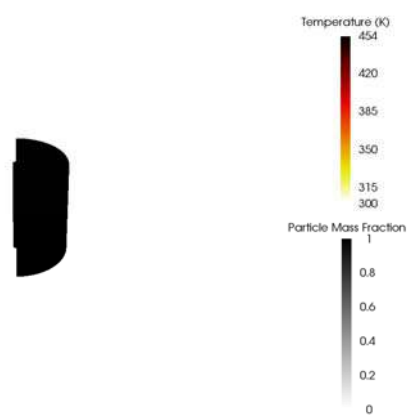

- Features:
  - 60 Million computational cells
  - 15 Million computational particles
  - r_{max} = 4.00m
  - t_{max} = 2.50ms
  - 5760 cores
  - 1 mode perturbation in the initial particle distribution

CCMT

14

Demonstration Problem: Predictions

Blast Wave Location (PM-1)

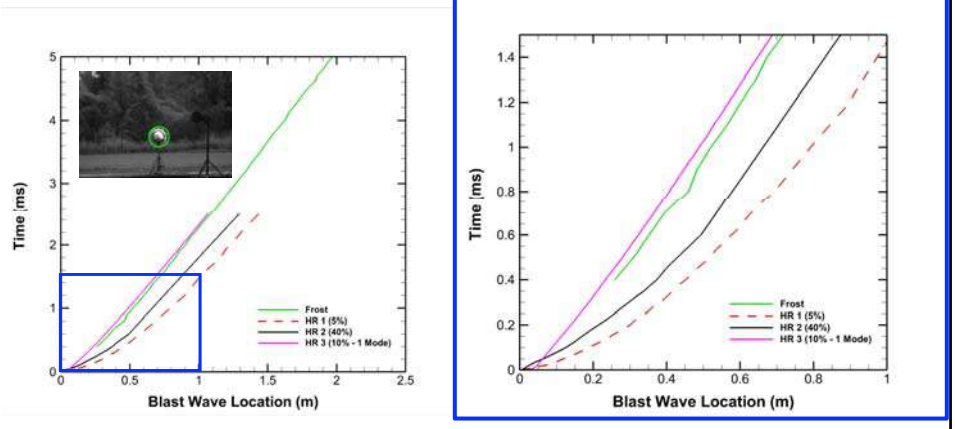

CCMT

15

Demonstration Problem: Predictions

Particle Front Location (PM-2)

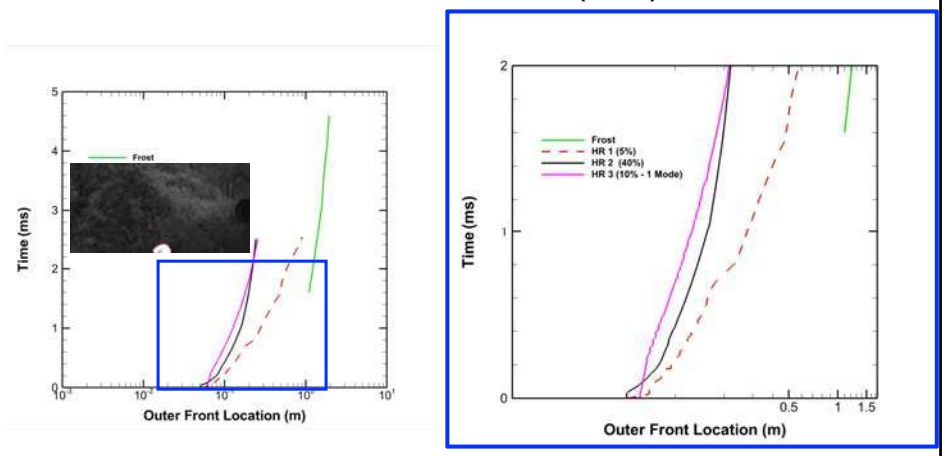

CCMT

16

Initial Condition Study

Uniform **1 Mode**

- The initial volume fraction distribution leaves an imprint in the fluid flow field
- The initial volume fraction distribution creates a channeling effect

CCMT

17

Amplification of Departure From Axisymmetry

1. Initial multi-modal **perturbation** in the particles

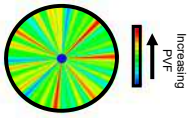

2. Metric of **departure** from **axisymmetry**

$$\zeta(t) = \frac{[Max(PVF(\theta)) - Min(PVF(\theta))](t)}{[Max(PVF(\theta)) - Min(PVF(\theta))](t=0)}$$

3. Which **initial disturbances amplify most** the departure?

4. High number of **simulations (2500)**

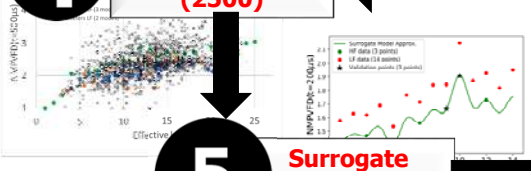

5. **Surrogate** Models

6. **Optimization**

➢ Giselle Fernandez, tomorrow 10:30 am

CCMT

18

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)

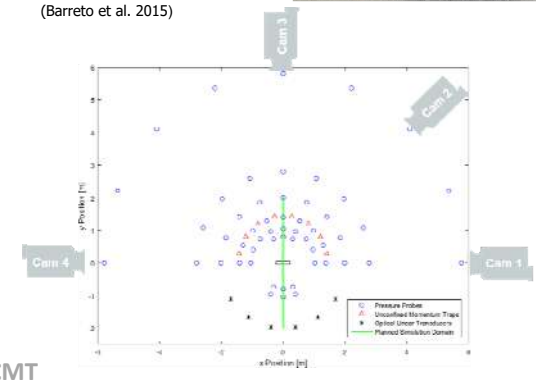

- Instrumentation:
  - 54 pressure probes
  - 8 Momentum traps
  - 4 high speed video cameras
  - 6 Linear optical transducers

CCMT

19

Eglin Blastpad Experimental Setup

(Barreto et al. 2015)

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

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

CCMT

19

Eglin Blastpad Simulation Setup

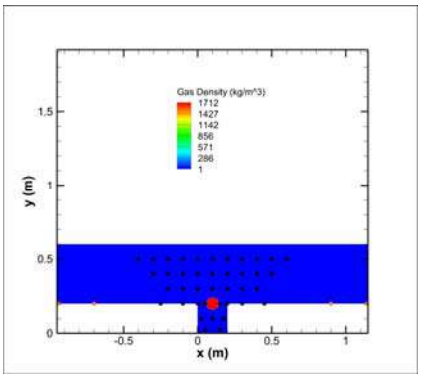

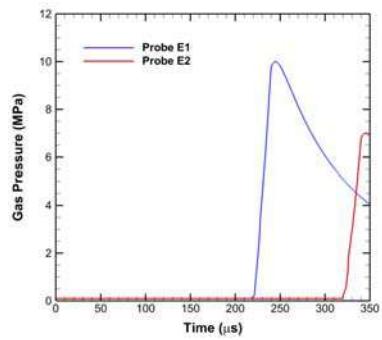

Note: Bare charge setup is shown, particles will be included in parallel runs

Preliminary Simulation Details:

- 1 million cells in two dimensions
- To be run on 256 cores using LLNL's Quartz machine
- Exhaust boundary is an outflow
- Other bottom boundaries are slipwalls
- 4 probes set in locations of experimental sensors
- 40 other probes set in near-charge region

CCMT

20

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 in progress
- ✓ Blastpad tests in progress

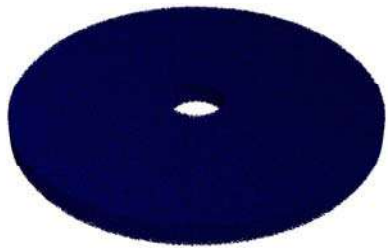

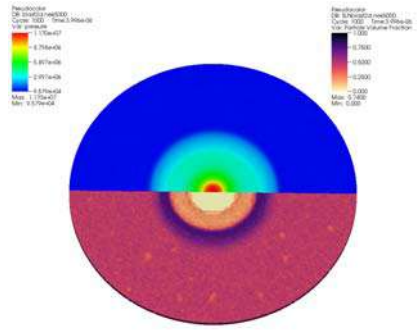

➢ David Zwick, 11:00 am

CMT-Hero 2

CCMT

21

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

CCMT

22

Eglin Microscale/Mesoscale Experiments

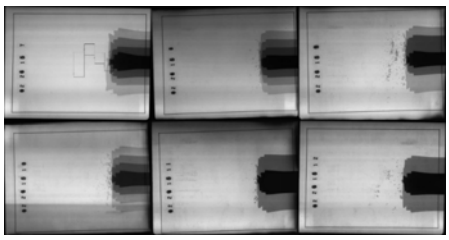

CCMT

23

Eglin Microscale Experiments

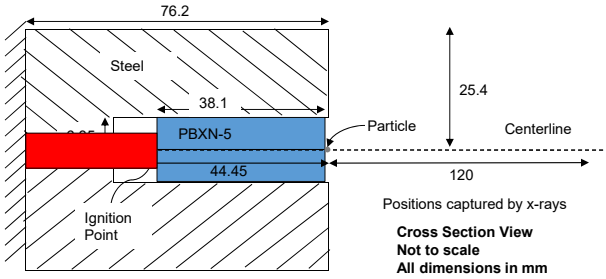

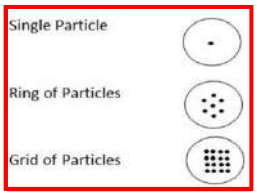

- Simulations implement the JWL equation of state and reactive burn initial condition for the explosive
- Plan for Uncertainty Quantification Analysis
  - Vary input parameters of particle density, particle diameter, and explosive mass for batch run (5 x 5 x 5 x 3 simulations) to construct surrogate model
  - Done for all variations of particle configuration
- Experiment has a vent in the back of the explosive where the detonator attaches
  - Simulations will account for the effects of this vent

CCMT

24

Mesoscale Simulation of Microscale Experiment

Geometric Features

Explosive Initial Conditions

- Lumped detonation
- Two Material Reactive Burn

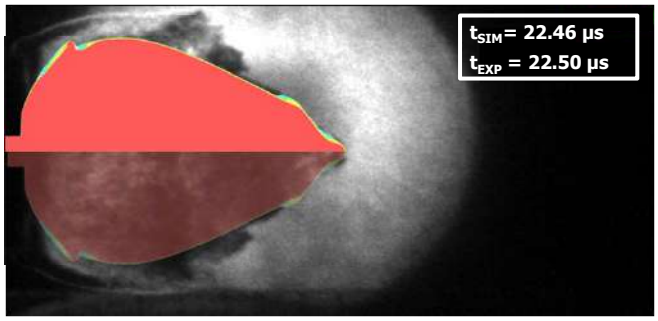

CCMT

25

Mesoscale Simulation of Microscale Experiment

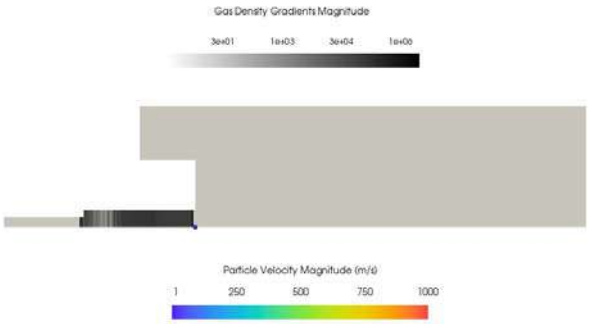

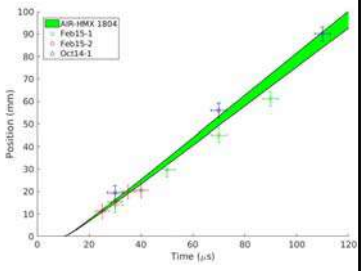

Validation Data Details:

- Drag models same as for Sandia's shock-tube simulation campaign
- Crosshairs show experimental data points with uncertainty
- The filled band outlines the max and min trajectories of the 25 particle
- Simulation results are time shifted based on Reactive Burn detonation times

CCMT

26

Mesoscale Simulation of Mesoscale Experiment

P-Rad Mesoscale Experiment (x-ray)

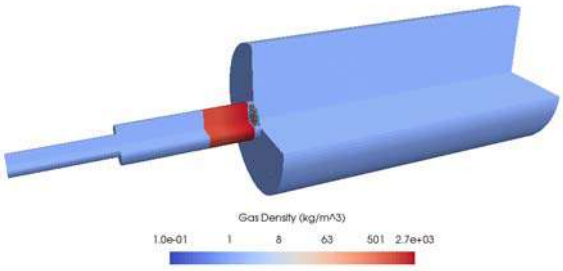

- Kyle Hughes, tomorrow 10:00 am

Simulation Details:

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

CCMT

27

UF UNIVERSITY of FLORIDA

CMT-nek Micro-Meso-Hero Timeline

		Y4-Q1	Y4-Q2	Y4-Q3	Y4-Q4	Y5-Q1	Y5-Q2	Y5-Q3	Y5-Q4
S1	**Micro Sims**								
	4-way coupling								
	Load Balancing								
	Compaction								
S2	**ASU Meso**								
	Outflow BC								
	Improved models								
S3	**Eglin Meso**								
	EVM tuning								
	Smart start								
	Real gas EOS								
	Reactive Burn								
S4	**CMT-Hero**								

- Goran Marjanovic
- David Zwick
- Josh Garno
- Rahul Koneru
- Fred Ouellet

CCMT

28

CCMT

Do you have any questions?

UF UNIVERSITY of FLORIDA

CCMT

CMT-nek:
A compressible multiphase DGSEM extension to nek5000

Jason F. Hackl
David Zwick

UF UNIVERSITY of FLORIDA

CMT-nek users and developers


Prof. Paul Fischer, UIUC


Jason Hackl, UF


David Zwick, UF


Tania Bannerjee, UF


Nguyen Tri Nguyen, UF

Nguyen joined December 2017

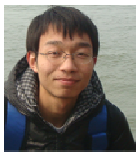
Li Lu, UIUC


Brad Durant, UF


Goran Marjanovic, UF


Rahul Koneru, UF


Fred Ouellet, UF


Keke Zhai, UF

CCMT

2

Accomplishments in Year 4

- Discontinuous Galerkin SEM of 2nd derivatives
  - Entropy viscosity for stabilizing shocked flow
- Particle 4-way coupling
- Compressible multiphase flows with shocks
- Transition to CMT-nek
  - Microscale: shock-sphere interactions
  - Mesoscale: ASU expansion
  - Macroscale: CMT-Hero 1: blast with particles
  - Detonation conditions (gas only)

CCMT 3

Outline

1. CMT-nek gas dynamics
   a. Discontinuous Galerkin method for gas flow solver
   b. Entropy viscosity method
   c. V&V of shock-capturing and microscale flow
2. CMT-nek multiphase
   a. Lagrangian particles and scaling
   b. Path to demonstration problem
3. Code development and future work
   a. Transition and timeline
   b. Risk mitigation
   c. Future work

CCMT 4

Discontinuous Galerkin Spectral Elements

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

flux = convective + diffusive
$\mathbf{H}_i = \mathbf{H}_i^c(\mathbf{U}) + \mathbf{H}_i^d(\mathbf{U}, \nabla \mathbf{U})$

Right-hand-side for fully explicit TVDRK3

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

Inviscid Euler equation fluxes $\mathbf{H}^c$

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

Artificial viscosity $\mathbf{H}^d$ + source terms

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

- Mass matrix $\mathbf{B}_{ik}$

$$\iiint_{-1}^{1} f(\mathbf{r}) J d\mathbf{r} \approx \mathbf{B}\mathbf{f},\ B_{ij} = \omega_i f(r_i) J \delta_{ij}$$

Numerical flux $\mathbf{H}^*(\mathbf{U}\text{-},\mathbf{U}+)$

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

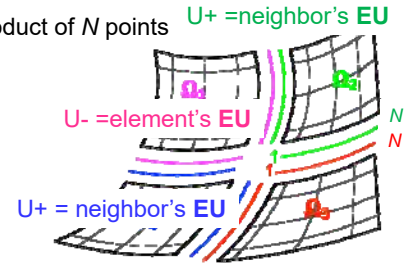

$$\{\{u\}\} \equiv \frac{1}{2}\left(u^- + u^+\right) = \frac{1}{2}\mathbf{QQ}^\top[\mathbf{Eu}] \qquad \mathbf{u}^+ = \left(\mathbf{QQ}^\top - \mathbf{I}\right)[\mathbf{Eu}]$$

CCMT

5

Discontinuous Galerkin Spectral Elements

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

fully explicit; time step limited by diffusion

$$D = \frac{\mu \Delta t}{h^2},\ D \ll 1$$

$$I_{vol}^{(c)} = \left[\mathcal{I}^\intercal \mathbf{D}_{M,k}^\intercal \mathbf{B}_M \mathcal{I}\left[\mathbf{r}_{k,\mathbf{x}_i}\right]\left[\mathbf{H}_i^c(\mathcal{I}\mathbf{U})\right]\right.$$

Volume: $O(N_{elem} M^4)$

$$I_{sfc}^{(c)} = \left[\mathbf{E}^\intercal \mathcal{I}_A^\intercal \mathbf{B}_{M,A}\left[\mathbf{H}_i^{c*}\left(\mathcal{I}\left[\mathbf{EU}\right]\right)\right]\left[\mathcal{I}_A n_i\right]\right.$$

Surface: $O(N_{elem} M^2)$

Interpolation to *M*=3(N-1)/2 GL points via matrix $[\mathcal{I}$ for polynomial dealiasing

isoparametric transformation metrics

Finite difference op **D** $\frac{\partial}{\partial x_i} = \frac{\partial r_k}{\partial x_i}\frac{\partial f}{\partial r_k} \approx \left[\mathbf{r}_{k,\mathbf{x}_i}\right]\left[\mathbf{D}_k \mathbf{f}\right]$

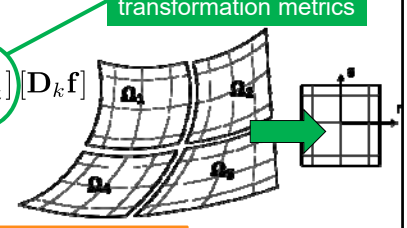

Ronquist and Patera (1987) *Intl. J. Numerical Methods Engrg.* **24**, 2273-2299
Fornberg (1998) *A Practical Guide to Pseudospectral Methods* Cambridge
Gottlieb & Shu (1998) *Math. Comp.* **67**:73-85
Deville, Fischer and Mund (2002) *Higher-Order Methods for Incomp. Flows* Cambridge

6

DGSEM: diffusive flux integrals

Quasi-linear flux Jacobian *A*
NOT weighted by volume fraction ϕ_g

$$H^d_{ij} = \mathscr{A}_{ijkl}\frac{\partial\left(U_k/\phi_g\right)}{\partial x_l}$$

i^{th} equation, flux in x_j

Weighted residual, integrate by parts → primal form → DGSEM ops

$$\int_{\Omega_e} v\nabla\cdot\mathbf{H}^d_i\,dV \to I_{KU} - I_{GU} + I_{G^\intercal U} \qquad \Gamma \equiv \bigcup_e \partial\Omega_e$$

$$I_{GU} = \int_\Gamma \{\{\mathscr{A}\,\nabla\mathbf{U}\}\}\cdot[\![v]\!]\,dA \to \mathbf{v}^\intercal\mathbf{E}^\intercal\mathbf{B}_A\left[\left(\mathbf{H}^d_i\cdot\mathbf{n}\right)^- - \{\{\mathbf{H}^d_i\cdot\mathbf{n}\}\}\right]$$

$$I_{G^\intercal U} = \int_\Gamma \{\{\mathscr{A}^\intercal\nabla v\}\}\cdot[\![\mathbf{U}]\!]\,dA \to \mathbf{v}^\intercal\mathbf{D}^\intercal_j\mathscr{A}_{ij}\mathbf{E}^\intercal\mathbf{B}_A\left[\mathbf{U}^- - \{\{\mathbf{U}\}\}\right][n_i]$$

$$I_{KU} = \int_{\Omega_e}(\nabla v)\cdot\mathbf{H}^d_i\,dV \to \mathbf{v}^\intercal\mathbf{D}^\intercal_k\,\mathbf{B}\left[\mathbf{r}_{k,\mathbf{x}_i}\right]\mathbf{H}^d_i$$

Baumann & Oden (1999) *Comput. Methods Appl. Math. Engrg.* **175**:311-341

Dolejší, V. (2004) *Intl. J. Numer. Meth. Fluids* **45**:1083-1106
Hartmann & Houston (2008) *J. Comp. Phys.* **227**:9670-9685

CCMT

7

Entropy viscosity method

Artificial stress tensor[1] (**NOT** Navier-Stokes) with AV that tracks the **entropy residual**
NOT weighted by volume fraction ϕ_g

Diffusive flux of…

Mass

$$\mathbf{H}^d_1 = -\kappa_s\nabla\rho$$

Momentum

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

$$\sigma_{ij} = \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}$$

Energy

$$\mathbf{H}^d_5 = -\left[\rho\,\nu_s\mathbf{u}:\underline{\sigma} + \kappa_s\,\mathrm{X}\left(\nabla(\rho e) + \frac{1}{2}|\mathbf{u}|^2\,\nabla\rho\right)\right]$$

TUNEABLE

due to **entropy viscosity**

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

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

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

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

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

R_s approx by finite diff between RK3 stages' s

$$s = \rho\,c_v\log\left(p\rho^{-\gamma}\right)$$

C^0 by average R_s at face points. Smooth via

$$S(R) = \frac{1}{12}\left(6R_{i,j,k} + R_{i+1,j,k} + R_{i-1,j,k} + \ldots\right)$$

[1]Guermond & Popov (2014) *J. Appl. Math.* **74**:284-305
Guermond, Pasquetti & Popov (2011) *J. Comp. Phys.* **230**:4248-4267

8

Convergence and validation

- Mass diffusivity $\kappa = P\,\nu$ regularizes solution

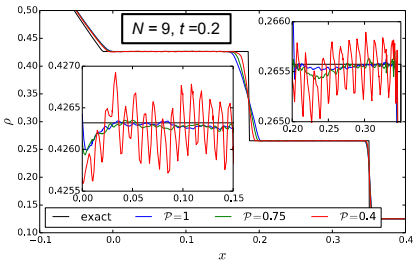

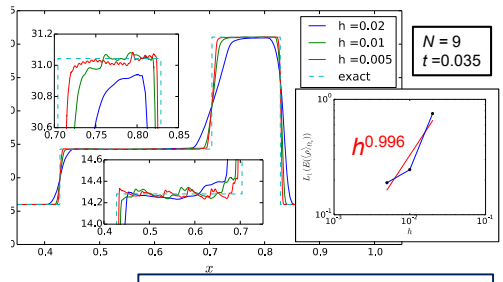

- ν_* reaches ν_{max} at shock wave

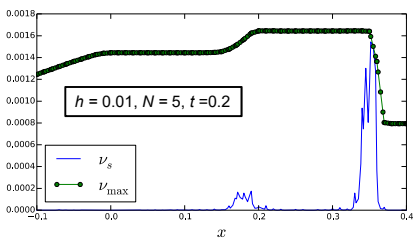

Calibrated P, c_E, c_{max}
Sod (1978) shock tube
Toro (2001) two-shock collision

	Sod (1978)	2-shock
c_{max}	0.3	1.0
c_E	40	85
P	0.75	0.75

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

9

Shock-sphere drag coefficient

- Mach 3 shock-sphere interaction, $N = 5$

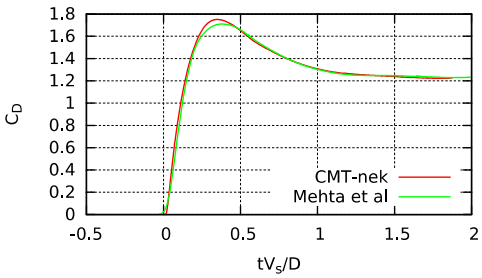

- Mach 1.22 shock-sphere, two resolutions

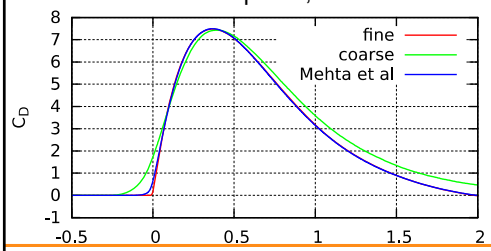

	Mach 1.22	Mach 3
c_{max}	0.3	0.75
c_E	40	85
P	0.75	0.75

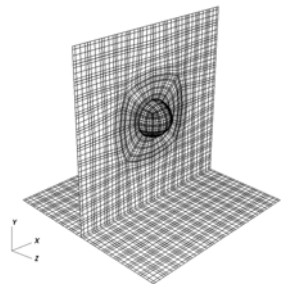

Coarse: 2048 elements, $N = 5$
Fine: 1 M elements, $N = 5$

Mehta, Jackson, Zhang & Balachandar (2016) *J. Appl. Phys.* **119**:104901
Hackl, Shringarpure, Koneru, Delchini & Balachandar (2018) in revision for *Computers & Fluids*

10

Shock-sphere interaction

Mach 1.22 | Mach 3

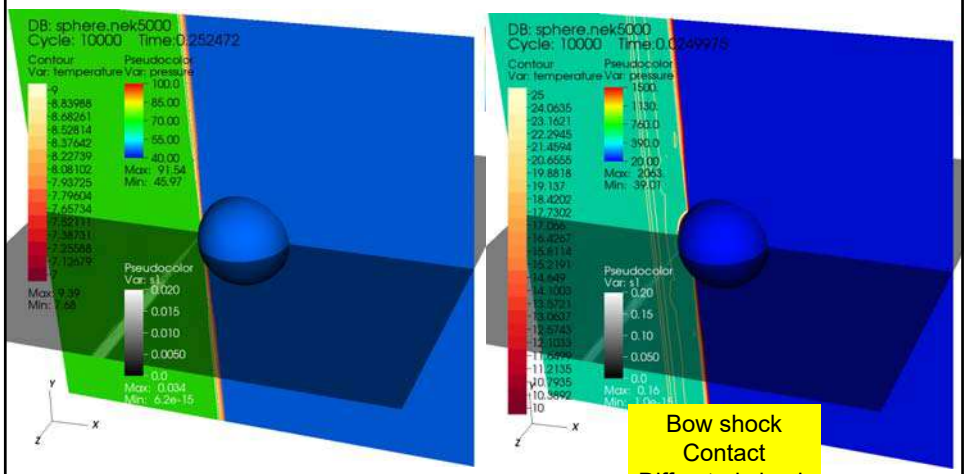

Captured features identified in experiments & sim

- Tanno *et al* (2003) *Shock Waves* **13**(3):191-200
- Bryson & Gross (1961) *J. Fluid Mech.*
- Mehta, Jackson, Zhang & Balachandar (2016) *J. Appl. Phys.* **119**:104901

CCMT

11

Microscale: Potential flow over spheres

Tests generalized Faxen theorem[1]

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

N_{elem}	30 K
DOF	65 M
Procs	8192

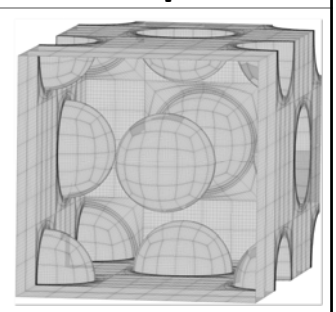

Mach

0.000 0.1809 0.3619 0.5428 0.7237

FCC 10% Volume fraction

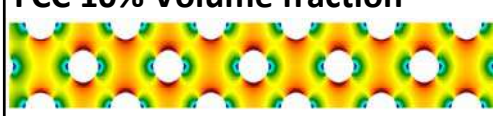

FCC 15% Volume fraction

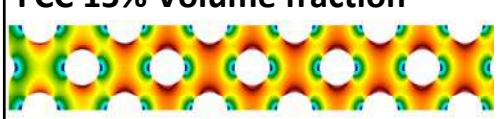

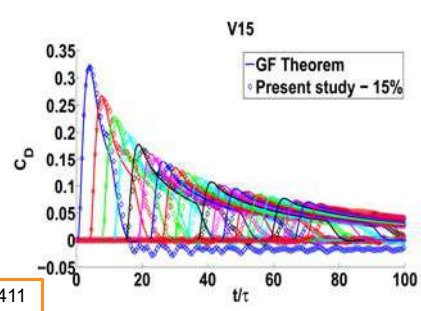

[1]Annamalai & Balachandar (2017) *J. Fluid. Mech.* **816**, 381-411
Marjanovic *et al* (2018) in preparation for *Phys. Fluids*

12

Ideal-gas Demonstration problem

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

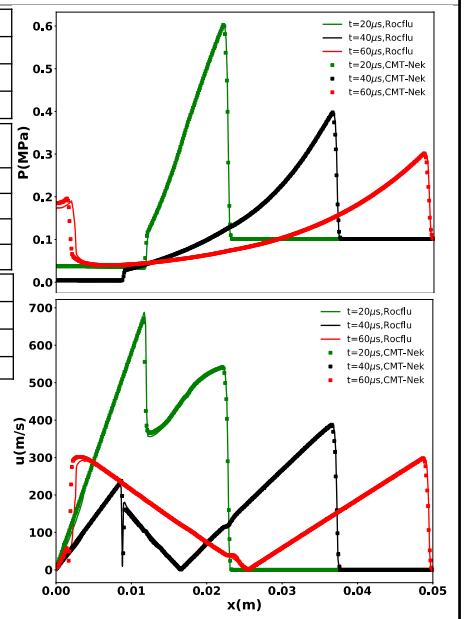

2D Verification-Gas only

Rocflu: 2D, ~ 2.5E6 cells, dr=154 μm

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

CCMT

13

CMT-Hero 1 (demo problem)

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

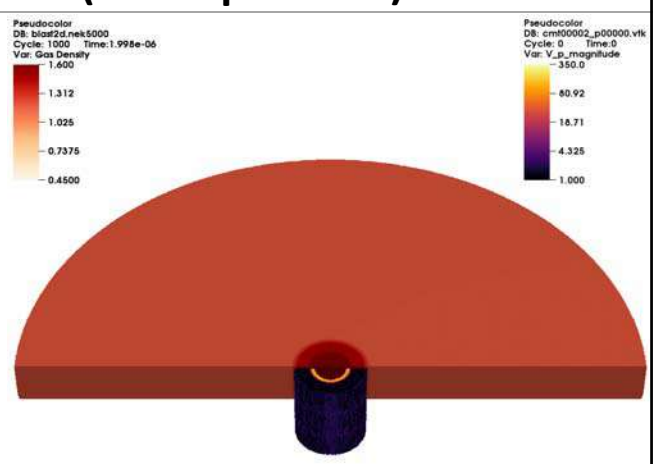

CCMT

14

Multiphase nek5000 and CMT-nek

- CMT-nek is built on single-phase nek5000
- While not fully resolved, Eulerian-Lagrangian method can simulate a more realistic number of particles
- Challenges related to particle-fluid coupling
  - Physics, numerics, computer science
- Outline:
  - Particle implementation in CMT-nek/nek5000
  - Current validation efforts
  - Multiphase summary

CCMT 15

Multiphase Capabilities

✓ 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 16

Multiphase Source Terms

Projection: Particles Effect Surrounding Grid

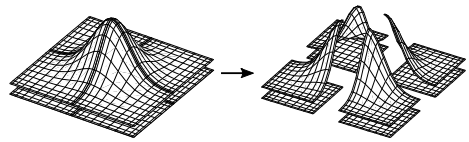

Particles & Fluid Grid

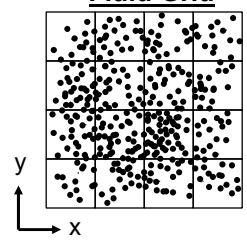

$$\mathbf{R} = \begin{bmatrix} 0 \\ \mathbf{f}_{pg} \\ \sigma_g : \nabla(\phi_p \mathbf{v}) + g_{pg} + q_{pg} \end{bmatrix}$$

Particle-fluid force coupling and energy contribution

Particle-fluid energy coupling

Eulerian particle velocity

CCMT

17

Governing Equations

For a single particle:

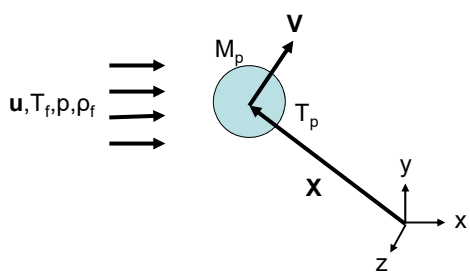

Position:

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

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

18

Inter-Particle Collision Model

Normal spring-damper soft-sphere collision model:

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

From thousands of particles…

… to millions of particles…

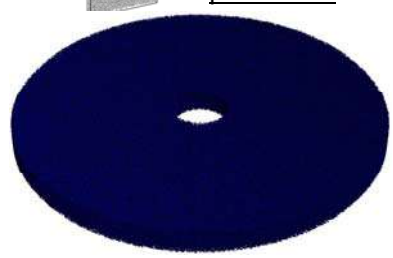

… and even billions…

Spring force:

- Overlapping particles repel each other

Damping force:

- Colliding particles dissipate energy

CCMT

19

Ghost Particles

If a particle is near a MPI rank edge, it will create a copy of itself called a ghost particle

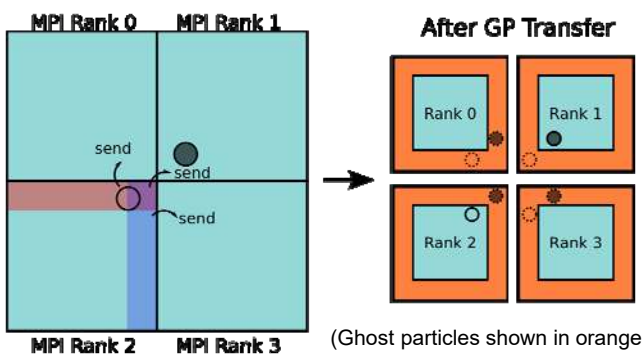

(Ghost particles shown in orange boxes)

Used for:

1. Particle-fluid projection
2. Particle-particle collisions

CCMT

20

Algorithmic Scaling

Ideal Setup:

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

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

Strong Scaling (s):

- 15 million grid points
- 8.4 million particles

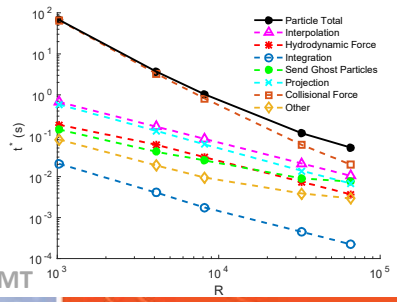

Weak Scaling (w):

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

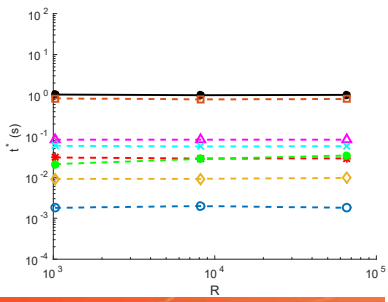

CCMT

21

Load Balancing on CMT-nek

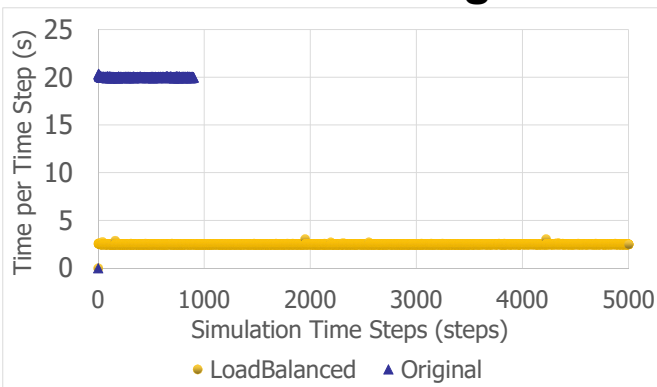

Vulcan

65,536 MPI ranks
16,384 nodes
4 cores per node
900,000 elements
1.125×10^9 particles
Grid size: 5x5x5
Rarefaction test
7.9x improvement in performance

- Time per time step:
  - 20.00 s for the original code
  - 2.52 s for the load balanced code
- 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

CCMT

22

Particle Verification

Coupling hierarchy:

- Fluid-particle (one-way)
  - One particle settling
- Particle-fluid (two-way)
  - Ghost particle creation
  - Eulerian vs. Lagrangian integration
- Particle-particle (four-way)
  - Single particle-wall collision
  - Nearest-neighbor count
  - Packed bed Voronoi volume distribution

CCMT 23

Fluidized Bed

- 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)	1000
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

Mesh and Boundary Conditions

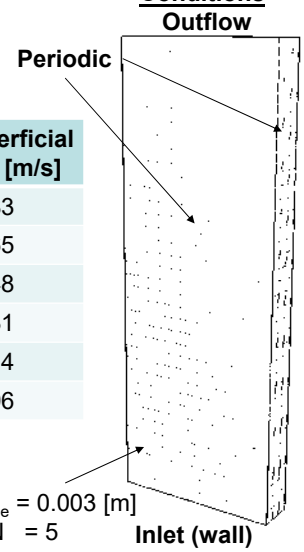

L_e = 0.003 [m]
N = 5

Simulations are performed with nek5000

CCMT 24

Fluidized Bed Results

Average pressure drop vs. inlet velocity

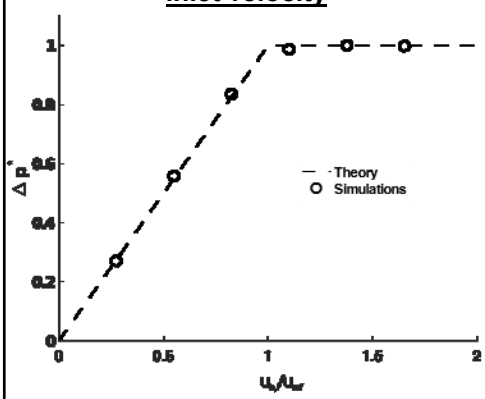

Particle volume fraction (Run 6)

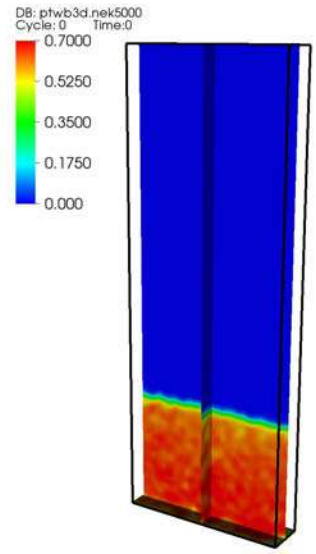

CCMT

25

Particles in Expansion Waves

Background:

- Problem consists of a vertical shock tube with particle bed in driver section
- Particle bed swells upwards due to rapid decompression when diaphragm bursts
- Results are compared with Cagnoli et al. (2002) experiments, which provide first version of ASU simulations at reduced scale
- Simulation is a smaller periodic slice of entire domain

Experimental Setup:

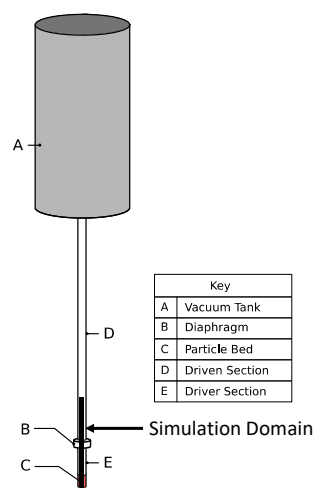

CCMT

26

Multiphase Refinement

- Grid refinement in multiphase flow is not well defined
- We consider two types of refinement:
  1. Fluid refinement (grid resolution)
  2. Particle refinement (coarse graining)

Coarse graining: The number of real particles that each computational particle represents

- Three cases of increased refinement:

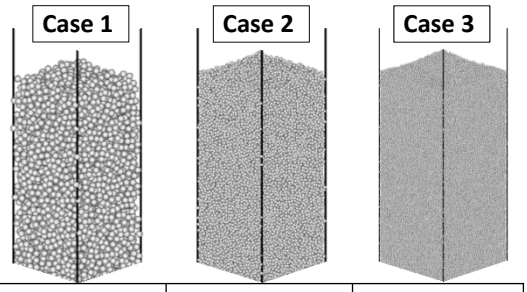

Coarse Graining	64	8	1
Particles	89,081	712,647	5,701,176
Grid Points	776,250	6,210,000	49,680,000
MPI Ranks	8k	65k	524k

CCMT

27

Simulation Setup

Geometry

	Experiment	Simulation
$L(cm)$	23	23
$D_H(cm)$	0.2	3.8
$h_0(cm)$	6.5	6.5

Particles

	Experiment	Simulation
$D_p(\mu m)$	38	38
$\rho_p(\mathrm{kg}/m^3)$	2,500	2,500
N_p	~2× 10^9	89k, 713k, 5701k
$k_c(kg/s^2)$	-	80, 40, 20
e_n	-	0.9

Grid

	Simulations
N_e	6k, 50k, 400k
N	5
$\Delta x(\mu m)$	167, 83, 42
$\delta_f/\Delta x$	2.6
δ_f/D_p	2.85

Initial Conditions

	Experiment	Simulation
$p_1(kPa)$	Laboratory	101
$T_1(K)$	Laboratory	293
$\Delta p\ (kPa)$	20 - 90	20 - 90
$T_4(K)$	Laboratory	293

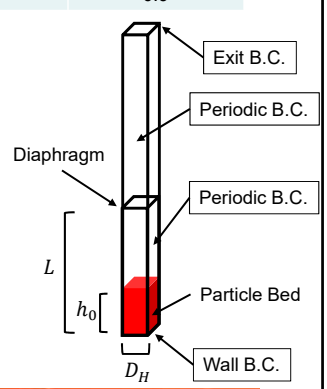

Simulations are performed with CMT-nek

CCMT

28

Particle Front Movement

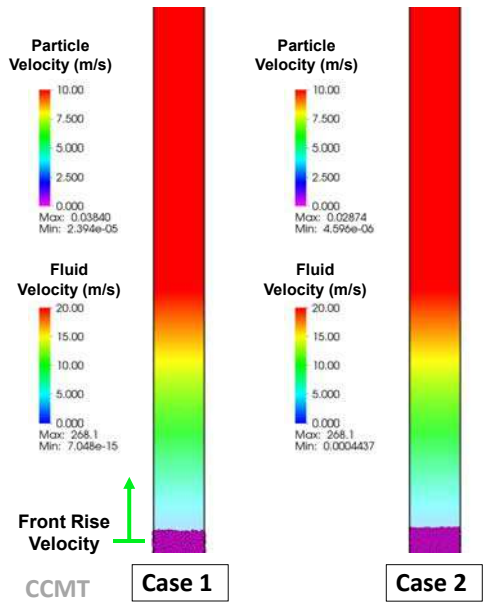

- Varying pressure difference across diaphragm
- Gidaspow drag force from fluidized bed community
- Measure initial velocities of bed height

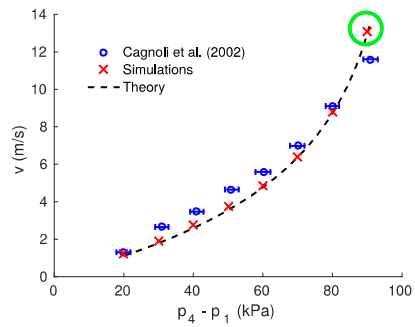

CCMT

29

Multiphase Summary

- Particles have been simulated across a wide range of physics
- Necessary features:
  - Efficiency (exascale systems)
  - Accuracy (physics)
  - Load-balancing (multiphase)
- Particles are capable of reproducing multiphase physics:
  - Fluidized bed
  - Particles in expansion waves

CCMT

30

UF UNIVERSITY of FLORIDA

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
- Load balancing accepted to ICS'18
- Automated regression testing via travis.cl
- Macroscale group in transition

CCMT

31

UF UNIVERSITY of FLORIDA

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		

CCMT

32

Transition plan: timeline

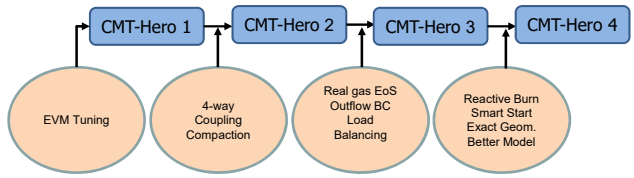

		Y4-Q1	Y4-Q2	Y4-Q3	Y4-Q4	Y5-Q1	Y5-Q2	Y5-Q3	Y5-Q4
S4a	**CMT-Hero 1**								
	EVM tuning								
S4b	**CMT-Hero 2**								
	Compaction								
	4-way coupling								
S4c	**CMT-Hero 3**								
	Real gas EOS								
	Outflow BC								
	Load Balancing								
S4d	**CMT-Hero 4**								
	Reactive burn								
	Smart start								
	Improved models								

CCMT

33

Transition plan: Risk Mitigation

Capability	Risk	Remediation / alternative
Artificial viscosity	c_E, c_{max} may not suffice for full duration of blast propagation	• Alternative AV to EVM • Lv&lhme (2015) EBDG solution limiter
Particle compaction	No compaction model. Post-shock oscillation at low ϕ_g may violate positivity	• Harris & Crighton • 2-fluid model at small t
Initialization and small time	Δt must start small; how should it grow with t?	• 2-fluid model at small t • Fluxes, RK scheme
Load balancing	Version control difficulties. ~1 element/task at t=0 for balanced particles	• 2-fluid ⟶ Lagr. PP. • Manual rebalancing at restarts
Boundary conditions	Need outflow BC allowing the blast wave to exit the domain.	• Sponge layers • Characteristic BC • Mesh stretching
Multi-species	Adding species for detonation products should be straightforward, but mixed EOS and artificial mass diffusion may interfere with multicomponent formulation.	• Replace GP14 with Navier-Stokes • 1-equation JWL-ideal-gas mixture • See "Artificial viscosity"
Reactive burn	Replace IC with prescribed field for elements initially lying in the explosive.	• Anticipate more severe small-time

AV: artificial viscosity
BC: boundary conditions
EBDG: Entropy bounded DG (*JCP* **295**:715-739)
JWL: Jones-Wilkins-Lee (1968) state equation
PP: point-particle approach
IC: initial conditions
EOS: equation of state
EVM: entropy viscosity method
GP14: Guermond-Popov (2014)

CCMT

34

Summary and Future Work

- CMT-nek is filling RocFlu's shoes and beyond
  - Compressible multiphase flow, scalable
  - Microscale, mesoscale and macroscale transition underway
- Plan to fill capability gaps
  - Risk mitigation plan identifies alternatives
  - Computer science integration
  - Load-balancing (multiphase)
- Future work (new capability)
  - Entropy stability
    - Entropy-variables splitting (Carpenter *et al* (2014) *SIAM J. Sci. Comput.* **36**)
    - Entropy-stable multiphase source terms
  - Mesh generation

CCMT

35

CCMT

Do you have any questions?

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

UF 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

Raj Rajagoplan
BE-SST

Ben Reynolds
Thermal Modeling

Chad Saunders
Symbolic Regression

CCMT

2

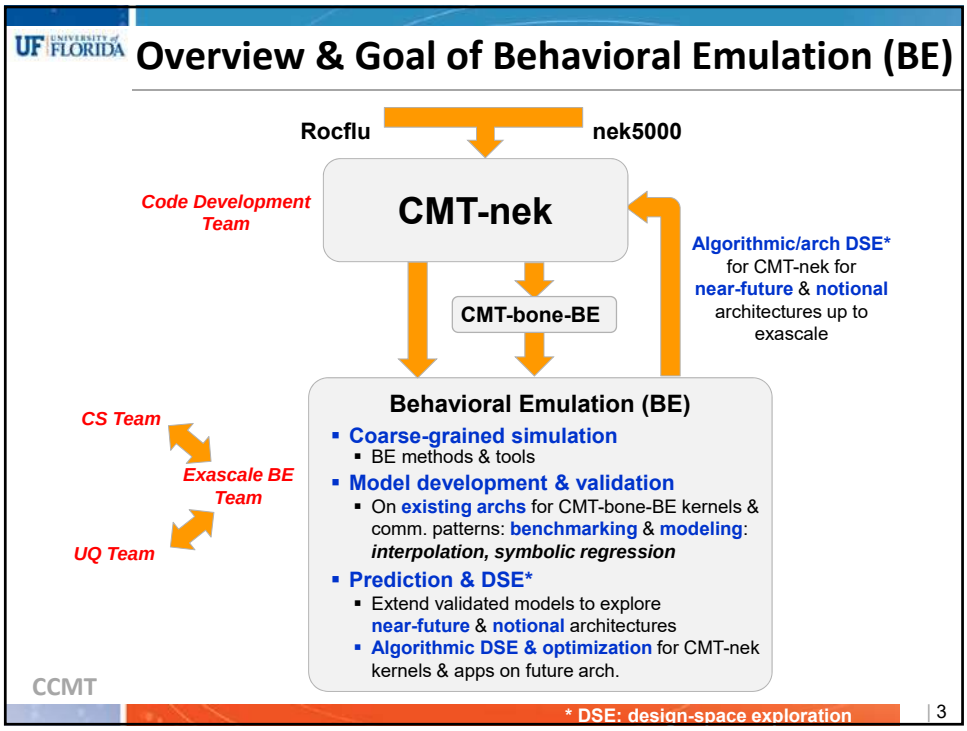

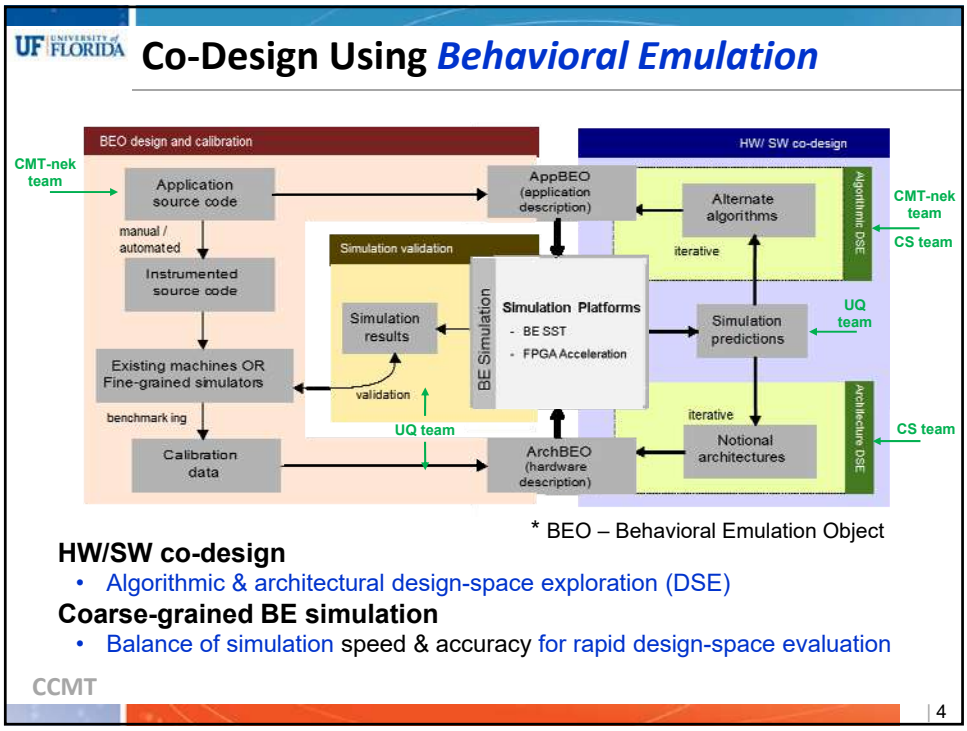

* BEO – Behavioral Emulation Object

HW/SW co-design

- Algorithmic & architectural design-space exploration (DSE)

Coarse-grained BE simulation

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

Outline

- Review & overview:
  - Behavioral Emulation (BE)

Research thrusts & achievements

  - BE-SST simulator & tools, experimental results
  - BE method enhancements
    - *Multi-fidelity surrogate model* for performance prediction
    - *Symbolic regression* for performance modeling
    - *Thermal* experiments & modeling
  - Acceleration using FPGAs
    - Speedup vs BE-SST
    - Scalability on multiple FPGAs
  - End-to-end case study - CMT-nek DSE using FPGAs
    - Design space exploration of CMT-nek
- Summary, conclusions, & future work

CCMT | 5

Key Achievements (Year 4)

- Large-scale experiments using BE-SST
  - Validation (100k+ MPI ranks) on Titan & Vulcan {< 9% error}
  - Predictive simulations approaching a million MPI ranks {Vulcan+ & Titan+}
  - Notional/near-future architecture - SUMMIT performance prediction
- Innovations in BE methods
  - UQ | Symbolic regression | Thermal modeling
- FPGA acceleration of BE simulation
  - Validation of dataflow, pipelined approaches (FEP, CP)
  - Scalability via multiple FPGAs
- End-to-end case study: CMT-nek DSE using FPGA
  - Algorithmic and parametric design space exploration
  - Achieved 10^6 times speedup over BE-SST

CCMT | 6

Outline

- Review & overview:
  - Behavioral Emulation (BE)

Research thrusts & achievements

  - BE-SST simulator & tools , experimental results
  - BE metho
    - *Multi-f*
    - *Symbo*
    - *Ther*
  - Accelerat
    - Speedu
    - Scalab
  - End-to-en
    - Design space exploration of CMT-nek
- Summary, conclusions, & future work

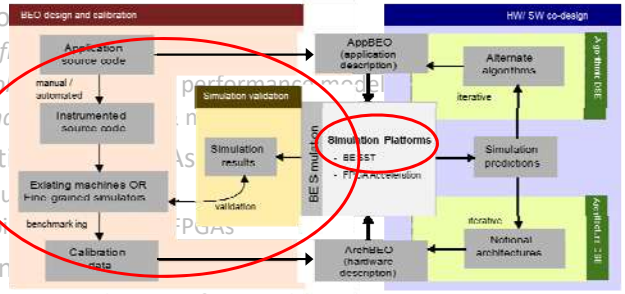

CCMT

BE-SST Simulator*

- Motivation: In-house BE simulator used to explore functionalities, but not very scalable or maintainable
- BE-SST: Developed by extending Structural Simulation Toolkit (SST)**
  - Framework for parallel simulations
  - Supported by developers & vendors
  - Flexibility in designing custom components

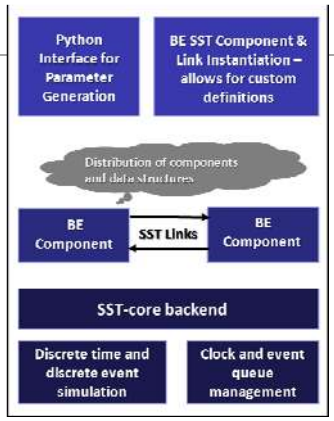

* A. Ramaswamy, N. Kumar, R. Rajagopalan, A. Neelakantan, H. Lam, Greg Stitt, "Scalable Behavioral Emulation of Extreme-Scale Systems Using Structural Simulation Toolkit", submitted to the 47th International Conference on Parallel Processing, Eugene, OR, August 13-16, 2018.

** From Sandia National Labs

SST Capabilities		BE Influences		BE-SST
• Parallel simulations • Discrete event simulation • Clock and event queues • Network Models • Component Models	+	• Software Definitions • Probabilistic Simulation • Abstract Network Definitions • Abstract Hardware Definitions	=	• Parallel discrete event simulation environment • Distributed component queues • Software Definitions • Probabilistic Simulation • Abstract Network Models • Abstract Component Models

CCMT

BE-SST Enhancements

- Communication model
  - Replaced static routing method with dynamic routing*
  - Inclusion of collectives such as MPI_ISEND, MPI_IRECV, MPI_WAITALL
- Performance model
  - Support for symbolic regression (in addition to interpolation API)
    - Symbolic regression with stochastic model
    - Improve accuracy (More details later)
- Leverage parallel capabilities of SST for BE-SST*
- Able to support trace-driven simulation**
  - For time-dependent simulations
  - To mimic CMT-nek with particles more accurately

* Details in Appendix
** Details in Sai Chenna's presentation on Friday

CCMT | 9

Outline

- Review & overview:
  - Behavioral Emulation (BE)

Research thrusts & achievements

  - BE-SST simulator & tools, experimental results
  - BE metho
    - *Multi-f*
    - *Symbo*
    - *The mo*
  - Accelerat
    - Speedu
    - Scalab
  - End-to-en
    - Design space exploration of CMT-nek
- Summary, conclusions, & future work

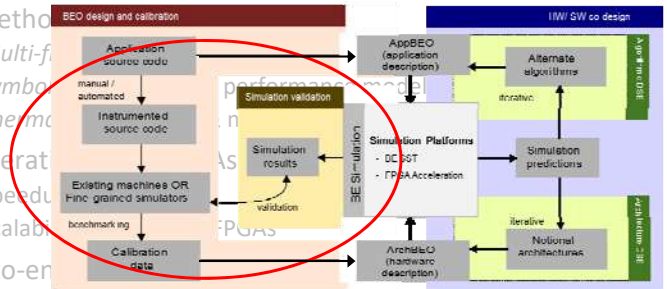

CCMT | 10

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

CCMT

11

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

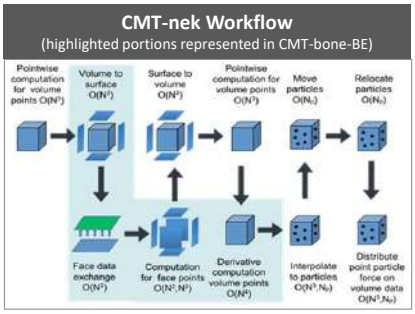

- Computation- & communication-intensive portions of CMT-nek workflow in C & MPI
  - Volume-to-surface data extraction
  - Face data exchange (with neighbors)
  - Derivative computation volume points
- Easier to modify for algorithmic DSE
- 1000 vs 10s of thousands lines of code

* design-space exploration

CCMT

12

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

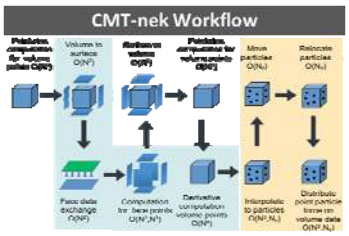

CCMT | 13

BE Simulations of Titan+

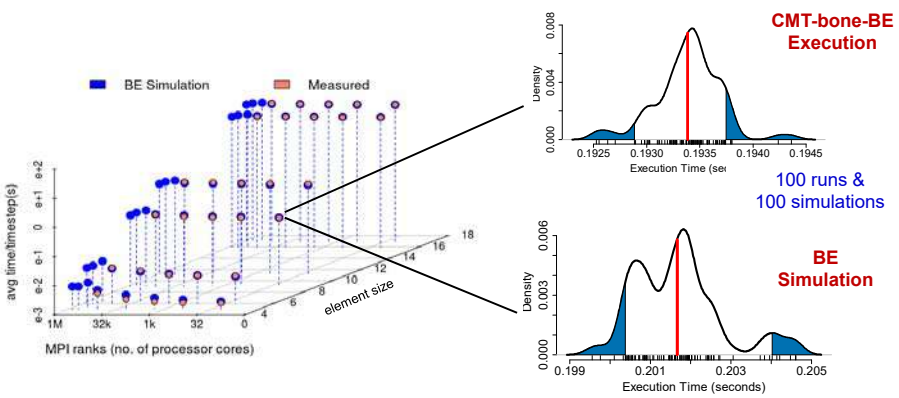

- Simulating bigger system than Titan (approaching 1 Million cores)
- Average % error between CMT-bone-BE simulation and execution time is 4%
- Maximum error is 17%

$$\%\,error = \frac{\frac{\sum_{i=1}^{N} t_{measured\,i}}{N} - \frac{\sum_{i=1}^{N} t_{simulated\,i}}{N}}{\frac{\sum_{i=1}^{N} t_{measured\,i}}{N}} \times 100$$

CCMT | 14

BE-SST Scalability

- Aiming for 1M+ cores; able to simulate 800K+
  - Memory bottleneck
- Uneven memory distribution across ranks while simulating
  - Bigger systems (e.g., 1 Million+ cores)
  - With more complex interconnects (e.g., 5D Torus of Vulcan)

*** GCM**
Graph Construction Memory: initial memory allocation before simulation

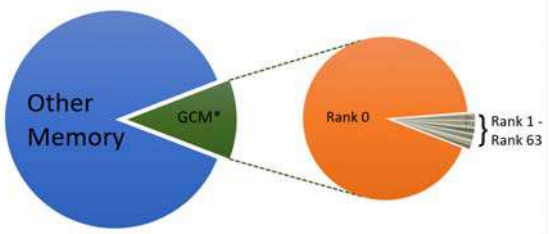

- Working with Sandia National Labs (Jeremy Wilke) to solve this issue

CCMT | 15

BE Simulations of Vulcan+

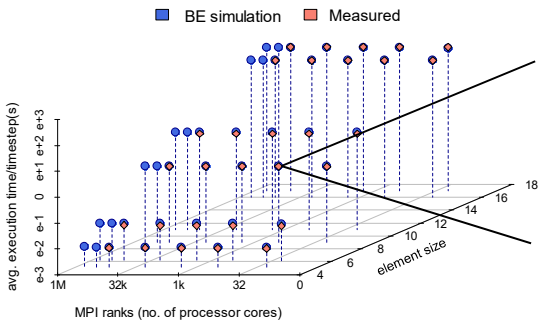

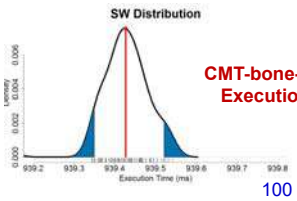

CMT-bone-BE Execution

100 runs & 100 simulations

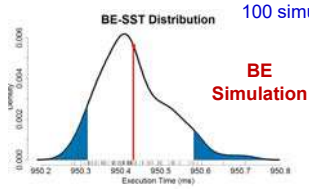

BE Simulation

- Simulating a bigger system than Vulcan (512k cores)
- Average % error between CMT-bone-BE simulation and execution time is 4%
- Maximum error is 9%

CCMT | 16

BE Simulations of Notional Architecture

- SUMMIT-like architecture
  - IBM POWER compute nodes
  - 44 cores/node
  - Non-blocking fat-tree n/w topology
- Initial methodology
  - Compute models of Vulcan (POWER architecture)
  - Fat tree architectural details and latency numbers obtained from Dr. Atchley Scott, ORNL
  - Predictions up to 4k nodes (180224 cores)
- Future work for notational exploration
  - Abstract Machine Modeling (AMM)*

* 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

CCMT | 17

SUMMIT-like Architecture

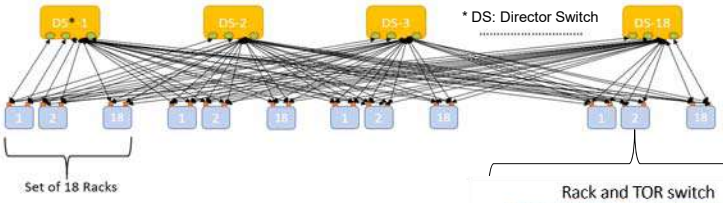

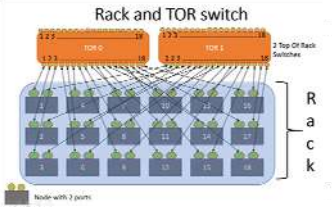

- System Performance: 5 to 10x Titan
- Number of nodes: ~4600
- Processor: IBM Power9
- Interconnect: Non-Blocking fat tree (dual-rail, single-plane)**
  - **Eighteen(18) 640-port** Director Switches
  - **256** Compute racks with **18** compute nodes/rack
  - Each rack has **two 36-port** Top-of-Rack (TOR) switches
  - Each node has **2 processors**; each processor has **22 cores** (44 cores/node)
  - Within rack: inter-node latency is **~ 0.7us** per hop
  - Outside of rack: additional **100us** per hop

CCMT

** Re-designed based on unpublished information obtained from Dr. Atchley Scott, ORNL

| 18

BE Simulations of SUMMIT-like Arch

Placeholder: anticipated performance results (at runtime)

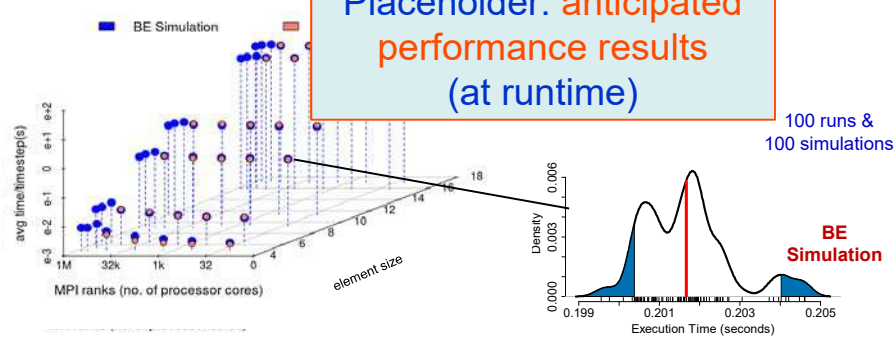

- Showcases capability of BE to perform simulations on notional architecture
- Cannot validate the result at this point
- Going forward
  - Refine models as new information is available (e.g., IBM Power 9 data)
  - Validate when access to SUMMIT is available

CCMT | 19

BE-SST Performance on Fat-tree Arch

Comparison GCM* between 5-D torus and non-blocking fat tree archs in BE-SST (Build-time)

* **GCM:** Graph Construction Memory: initial memory allocation before simulation

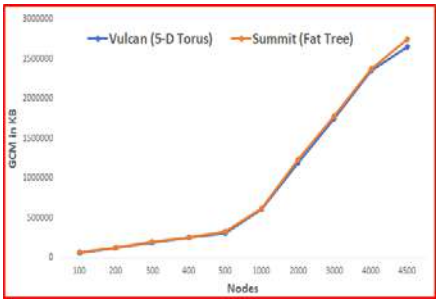

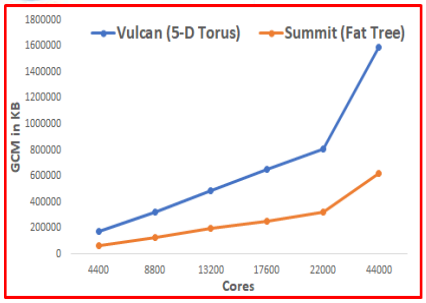

- In simulating equal # of ***nodes*** in BE-SST, GCM* for fat-tree configuration is approx. the ***same*** as the GCM* of 5-D torus
- In simulating equal # of ***cores*** in BE-SST, fat-tree configuration consumes only ***~40%*** of GCM* consumed by 5-D torus configuration

CCMT | 20

Summary: BE Methods, Simulator, & Experiments

- Large-scale experiments using BE-SST
  - Validation (100k+ MPI ranks) on Titan & Vulcan {< 9% error}
  - Predictive simulations approaching a million MPI ranks {Vulcan+ & Titan+}
  - Near-future architecture - SUMMIT performance prediction
- Going forward
  - BE-SST enhancements (e.g., scaling toward exascale)
  - BE-SST studies on new-existing, near-future, or abstract archs
- BE & BE-SST general accessibility & usability
  - Continual documentation:
    - BE-SST: architecture, code, configuration, usage
  - Planned availability
    - BE-SST code
    - Use cases: with complete build & execution walkthrough, expected results, assumptions/restrictions

CCMT | 21

Outline

- Review & overview:
  - Behavioral Emulation (BE)

Research thrusts & achievements

  - BE-SST simulator & tools, experimental results
  - BE method enhancements
    - *Multi-fidelity surrogate model* for performance prediction
    - *Symboli*
    - *Therma*
  - Acceleratio
    - Speedu
    - Scalabili
  - End-to-en
    - Design s
- Summary, c

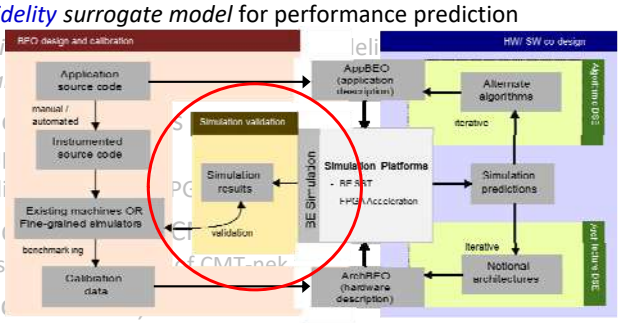

CCMT | 22

Behavioral Emulation with UQ*

w/ UB

Goal: Improve BE models via validation and uncertainty estimation

- Reduce computational budget by fitting BE simulation to CMT-nek using Multi-Fidelity Surrogate (MFS)
- Extrapolation of CMT-nek towards large-scale runs using BE and MFS

Corrected BE simulation against CMT-nek

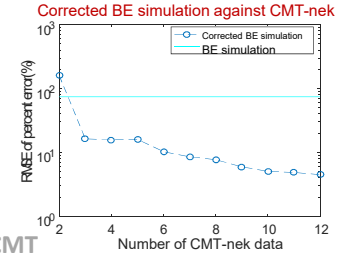

Accuracy of corrected BE sim.

- At 10 left-out CMT-nek test points
- Overall RMSE < 10% with 7 or more CMT-nek data
- Reducing computation budget
  - Supplementing high-fidelity (CMT-nek) with low-fid. (BE simulation) runs

Y. Zhang, A. Neelakantan, N. Kumar, C. Park, R. Haftka, N. H. Kim, and H. Lam, "Multi-fidelity Surrogate Modeling for Application/Architecture Co-design", Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS17), Denver, CO, Nov. 13, 2017.

*See Aravind's poster for more details

CCMT

23

Outline

- Review & overview:
  - Behavioral Emulation (BE)

Research thrusts & achievements

- BE-SST simulator & tools, experimental results
- BE method enhancements
  - *Multi-fidelity surrogate model* for performance prediction
  - *Symbolic regression* for performance modeling
  - T
- Acce
  - S
  - S
- End-
  - D
- Summa

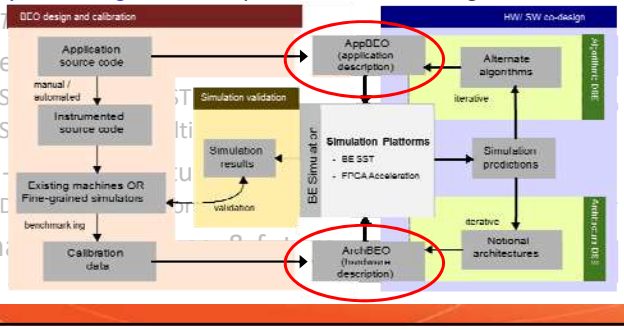

CCMT

24

Symbolic Regression: Motivation

- Motivation:
  - Multi-parameter performance modeling is challenging
  - Requires accuracy, computational efficiency, simple implementation
- Our previous approaches:
  - N-dimensional interpolation:
    - Computationally expensive
    - Difficult to implement, needs customization for each example
    - Insufficient accuracy
  - Multi-variate 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

CCMT 25

Symbolic Regression

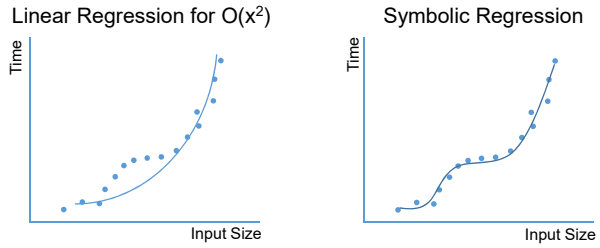

- Advantages:
  - Captures machine-specific performance behaviors
  - Enables tradeoffs between accuracy and computational complexity
    - Can maximize performance for a given error constraint
  - Requires no prior knowledge of the kernel
- Approach: custom genetic-programming tool
  - Evolves “trees” from mathematical primitives
  - Multiple optimization goals: minimize error, minimize error given tree size constraint, minimize tree size given error constraint, etc.

CCMT 26

Symbolic Regression: Results

CMT-nek particle solver kernel	Linear Regression	Symbolic Regression	Error Improvement
	MAPE	MAPE	$MAPE_{Linear}/MAPE_{symbolic}$
Compute1	3.5%	1.5%	2.4x
update_particle_location	10.8%	1.6%	6.8x
particles_in_nid	76.4%	31.1%	2.5x
interp_props_part_location	902.6%	10.0%	90.7x
upd_vel_and_pos_rk3_stage1	399.0%	9.0%	44.2x
upd_vel_and_pos_rk3_stage3	19.4%	9.7%	2.0x
upd_vel_and_pos_rk3_allstage	27.3%	6.8%	4.0x
upd_vel_and_pos_bdf	20.0%	4.9%	4.1x
red_interp	4.7%	4.3%	1.1x
usr_particles_forces	6.3%	2.1%	2.9x
tri_interp	1.4%	0.7%	1.9x
Average	133.8%	7.4%	14.8x

CCMT

27

Performance Modeling: Overhead Distribution

- Problem: symbolic regression provides deterministic value
  - Corresponding performance models do not include variance
- Solution: create separate models for minimum time and overhead
  - Execution time = minimum time + overhead time
  - Variation in overhead samples is modelled by a new parameter "rank"

$$Total\ Execution\ time = f(nx1, nelt, \alpha) + g(nx1, nelt, \alpha, rank)$$

$$3.68nx1^4\alpha + 1.71E^{-5} \qquad 3.54E^{-6}(e^{rank}\alpha nelt) + 2.871E^{-6}$$

$nx1$ = element-size; α = particles/gridpoint; $nelt$ = elements-per-processor; rank = randomized parameter

CCMT

28

Outline

- Review & overview:
  - Behavioral Emulation (BE)

Research thrusts & achievements

  - BE-SST simulator & tools, experimental results
  - BE method enhancements
    - *Multi-fidelity surrogate model* for performance prediction
    - *Symbolic regression* for performance modeling
    - *Thermal* experiments & modeling
  - Acce
    - S
    - S
  - End-
    - D
- Summa

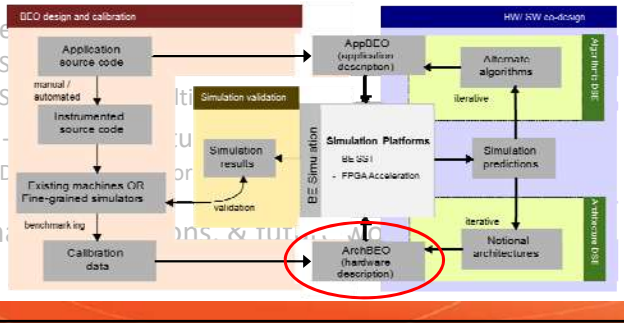

CCMT

29

Thermal Benchmarking & Modeling

- Previously modeled energy/power using physical and simulation measurements
- Current focus: thermal modeling with IR images of KNL
  - Goal: minimize maximum temperature via load balancing
- More challenging than anticipated
  - CPU fans block thermal information of CPU
  - Tried custom fan setup, destroyed several processors
- Modified KNL for IR by cutting hole in case:

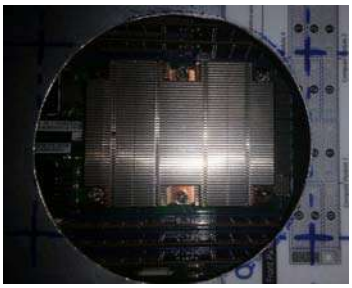

CCMT

30

Thermal Benchmarking & Modeling

Nek-5000 on KNL (64 cores)

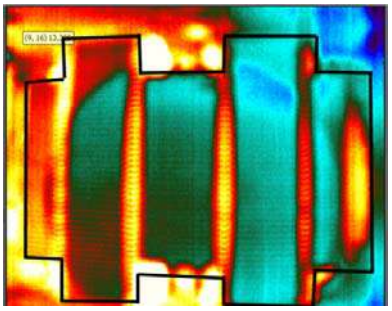

Nek-5000 on KNL (1 core)

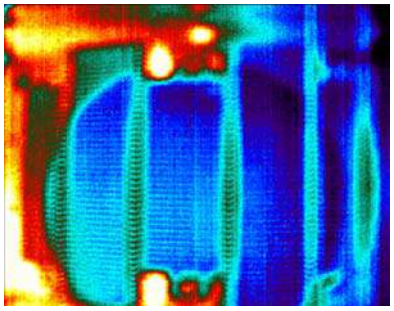

- Observations:
  - Different regions of the heat sink show widely different temperatures
  - Several columns show higher temperatures than others
    - Does not necessarily mean that processor is hotter at these points
  - Minimizing the maximum temperature *in the image* may not be most appropriate optimization goal
    - Need to weigh each region of the heat sink differently

CCMT

31

Thermal Benchmarking Methodology

- Map functionality onto specific cores using OpenMPI
- Identity 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

32 cores (Cores 0 to 31)

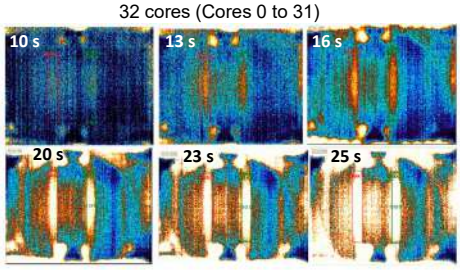

32 cores (Even cores: 0, 2, 4, …, 62)

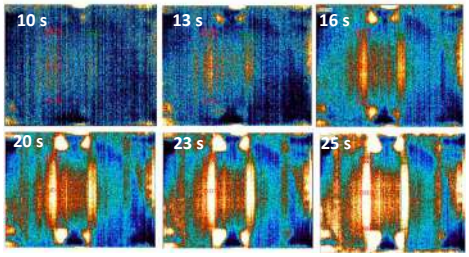

CCMT

32

Outline

- Review & overview:
  - Beha
- Research
  - BE-SS
  - BE m
    - M
    - Sy
    - Th
  - Acceleration using FPGAs
    - Speedup vs BE-SST
    - Scalability on multiple FPGAs
  - End-to-end case study - CMT-nek DSE using FPGAs
    - Design space exploration of CMT-nek
- Summary, conclusions, & future work

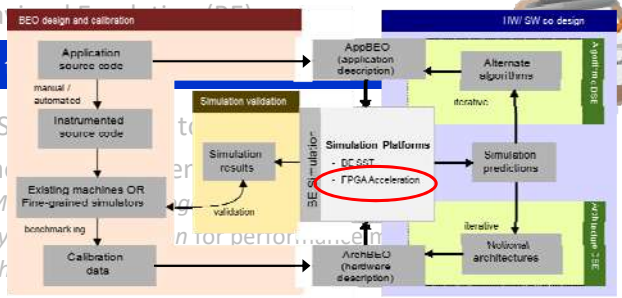

CCMT

33

FPGA-Accelerated Design-Space Exploration

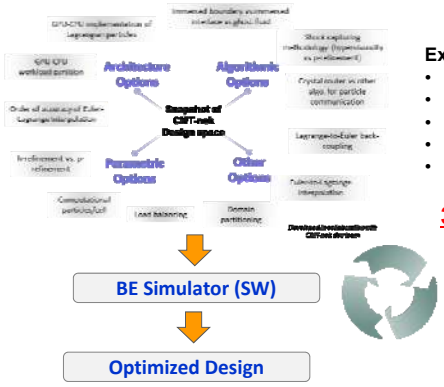

Example design space:

- **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
- See Carlo Pascoe's poster: FPGA-pipelined Simulations for CMT-nek

CCMT

34

Fully-expanded Pipeline (FEP)

1. Construct Data Flow Graph (DFG) from simulation configuration

- AppBEO+ArchBEO define instructions and operand/output dependencies
- Instructions map to vertices and dependencies map to edges in DFG
- Various opportunities for graph-level optimizations

1. Extracting DFG from BE simulation configuration

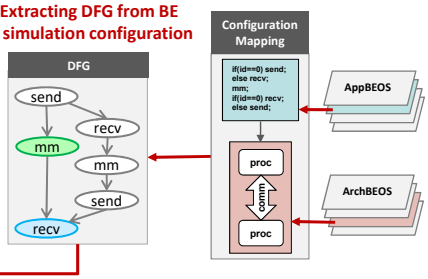

2. Mapping DFG to FPGA Pipeline

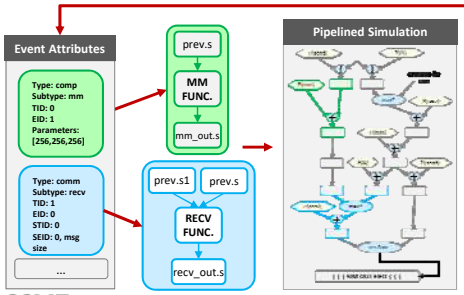

2. Map DFG to pipeline circuit

- Vertex attributes define operations and instantiate dedicated HW
- Edge attributes (e.g., src/dst) instantiate pipeline register between src/dst pair
- Various opportunities for circuit-level optimizations

Because each instruction (from sim) mapped to independent HW (no resource sharing), each vertex able to *start next sim 1 cycle after current sim*

CCMT

35

Collapsed Pipeline (CP)

1. Construct Data Flow Graph (DFG) from simulation configuration

- AppBEO+ArchBEO define instructions and operand/output dependencies
- Instructions map to vertices and dependencies map to edges in DFG
- Partition into linear subgraphs and generate dependency lists

1. Extracting DFG, then identify subgraphs & dependencies

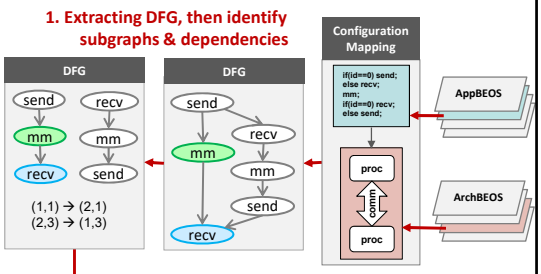

2. Mapping DFG to FPGA Pipeline

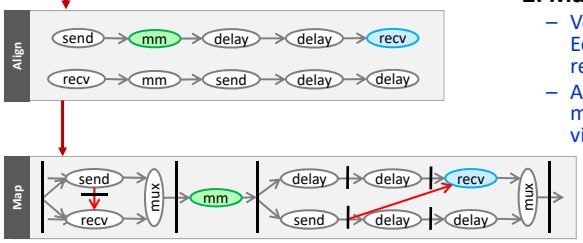

2. Map DFG to pipeline circuit

- Vertex attributes define operations and Edge attributes instantiate pipeline register between src/dst pairs
- Align subgraph traces such that cost is minimized and no dependencies are violated

Because each subgraph instruction mapped to independent HW, each vertex able to start next subgraph 1 cycle after current subgraph. *All subgraphs must complete before sim can complete*

36

UF UNIVERSITY of FLORIDA

Fully-Expanded & Collapsed Pipeline Tradeoffs

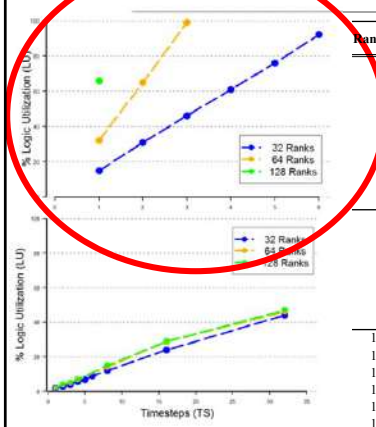

Ranks	TS	Num. of Events	% LU	Latency (cycles)	Hardware MSPS†	Hardware GEPS‡	BE-SST KEPS‡	Hardware Speedup
32	1	1,344	15 / 2	64 / 278	300 / 10.5	403 / 14.1	4.4	$92x10^6$ / $3x10^6$
32	2	2,688	31 / 3	118 / 512	300 / 10.5	806 / 28.2	7.8	$103x10^6$ / $4x10^6$
32	3	4,032	46 / 4	172 / 746	300 / 10.5	1,210 / 42.3	10.6	$114x10^6$ / $4x10^6$
32	4	5,376	61 / 6	226 / 980	300 / 10.5	1,613 / 56.4	12.9	$125x10^6$ / $4x10^6$
32	5	6,720	76 / 7	280 / 1,214	300 / 10.5	2,016 / 70.6	14.8	$136x10^6$ / $5x10^6$
32	6	8,064	92 / 9	334 / 1,488	300 / 10.5	2,419 / 84.7	16.5	$147x10^6$ / $5x10^6$
32	8	10,752	– / 12	– / 1,916	– / 10.5	– / 113	19.1	– / $6x10^6$
32	16	21,504	– / 24	– / 3,788	– / 10.5	– / 226	24.2	– / $9x10^6$
32	32	43,008	– / 44	– / 7,532	– / 10.5	– / 452	28.9	– / $16x10^6$
64	1	2,880	32 / 2	65 / 394	300 / 5.23	864 / 15.1	7.7	$112x10^6$ / $2x10^6$
64	2	5,760	65 / 4	119 / 712	300 / 5.23	1,728 / 30.1	12.6	$137x10^6$ / $2x10^6$
64	3	8,640	99 / 5	173 / 1,030	300 / 5.23	2,592 / 45.2	16.2	$160x10^6$ / $3x10^6$
64	4	11,520	– / 7	– / 1,348	– / 5.23	– / 60.2	17.9	– / $3x10^6$
64	8	23,040	– / 14	– / 2,620	– / 5.23	– / 121	23.2	– / $5x10^6$
64	16	46,080	– / 29	– / 5,164	– / 5.23	– / 241	26.9	– / $9x10^6$
64	32	92,160	– / 46	– / 10,252	– / 5.23	– / 482	29.1	– / $17x10^6$
128	1	5,952	66 / 2	66 / 458	300 / 2.62	1,786 / 15.6	10.8	$165x10^6$ / $1x10^6$
128	2	11,904	– / 4	– / 776	– / 2.62	– / 31.2	14.9	– / $2x10^6$
128	3	17,856	– / 5	– / 1,094	– / 2.62	– / 46.8	17.1	– / $3x10^6$
128	4	23,808	– / 7	– / 1,412	– / 2.62	– / 62.4	18.4	– / $3x10^6$
128	8	47,616	– / 15	– / 2,684	– / 2.62	– / 125	20.8	– / $6x10^6$
128	16	95,232	– / 30	– / 5,228	– / 2.62	– / 250	22.3	– / $11x10^6$
128	32	190,464	– / 47	– / 10,316	– / 2.62	– / 499	22.8	– / $22x10^6$

†Mega-Simulations-Per-Second, ‡Giga/Kilo-Events-Per-Second, '–' indicates configuration unable to fit on a single FPGA

Fully-Expanded Pipeline

Advantages:

- Superior performance in terms of simulation **throughput** and **latency**
  - 280 - 320 MHz implies 280 - 320 million simulations per second independent of simulated MPI ranks

Limitations:

- Resources scale linearly with both MPI Ranks and number of timesteps
- Scaling across multiple FPGAs expected to be ineffective when considering exascale simulations

CCMT

37

UF UNIVERSITY of FLORIDA

Fully-Expanded & Collapsed Pipeline Tradeoffs

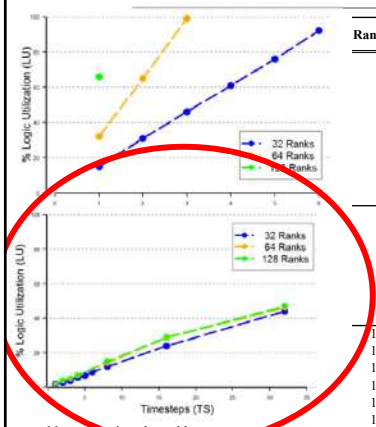

Ranks	TS	Num. of Events	% LU	Latency (cycles)	Hardware MSPS†	Hardware GEPS‡	BE-SST KEPS‡	Hardware Speedup
32	1	1,344	15 / 2	64 / 278	300 / 10.5	403 / 14.1	4.4	$92x10^6$ / $3x10^6$
32	2	2,688	31 / 3	118 / 512	300 / 10.5	806 / 28.2	7.8	$103x10^6$ / $4x10^6$
32	3	4,032	46 / 4	172 / 746	300 / 10.5	1,210 / 42.3	10.6	$114x10^6$ / $4x10^6$
32	4	5,376	61 / 6	226 / 980	300 / 10.5	1,613 / 56.4	12.9	$125x10^6$ / $4x10^6$
32	5	6,720	76 / 7	280 / 1,214	300 / 10.5	2,016 / 70.6	14.8	$136x10^6$ / $5x10^6$
32	6	8,064	92 / 9	334 / 1,488	300 / 10.5	2,419 / 84.7	16.5	$147x10^6$ / $5x10^6$
32	8	10,752	– / 12	– / 1,916	– / 10.5	– / 113	19.1	– / $6x10^6$
32	16	21,504	– / 24	– / 3,788	– / 10.5	– / 226	24.2	– / $9x10^6$
32	32	43,008	– / 44	– / 7,532	– / 10.5	– / 452	28.9	– / $16x10^6$
64	1	2,880	32 / 2	65 / 394	300 / 5.23	864 / 15.1	7.7	$112x10^6$ / $2x10^6$
64	2	5,760	65 / 4	119 / 712	300 / 5.23	1,728 / 30.1	12.6	$137x10^6$ / $2x10^6$
64	3	8,640	99 / 5	173 / 1,030	300 / 5.23	2,592 / 45.2	16.2	$160x10^6$ / $3x10^6$
64	4	11,520	– / 7	– / 1,348	– / 5.23	– / 60.2	17.9	– / $3x10^6$
64	8	23,040	– / 14	– / 2,620	– / 5.23	– / 121	23.2	– / $5x10^6$
64	16	46,080	– / 29	– / 5,164	– / 5.23	– / 241	26.9	– / $9x10^6$
64	32	92,160	– / 46	– / 10,252	– / 5.23	– / 482	29.1	– / $17x10^6$
128	1	5,952	66 / 2	66 / 458	300 / 2.62	1,786 / 15.6	10.8	$165x10^6$ / $1x10^6$
128	2	11,904	– / 4	– / 776	– / 2.62	– / 31.2	14.9	– / $2x10^6$
128	3	17,856	– / 5	– / 1,094	– / 2.62	– / 46.8	17.1	– / $3x10^6$
128	4	23,808	– / 7	– / 1,412	– / 2.62	– / 62.4	18.4	– / $3x10^6$
128	8	47,616	– / 15	– / 2,684	– / 2.62	– / 125	20.8	– / $6x10^6$
128	16	95,232	– / 30	– / 5,228	– / 2.62	– / 250	22.3	– / $11x10^6$
128	32	190,464	– / 47	– / 10,316	– / 2.62	– / 499	22.8	– / $22x10^6$

†Mega-Simulations-Per-Second, ‡Giga/Kilo-Events-Per-Second, '–' indicates configuration unable to fit on a single FPGA

Collapsed Pipeline

Advantages:

- Resources scale linearly with timesteps, but sublinearly with MPI Ranks
  - allows for significantly more timesteps due to its much lower base utilization
- Better scaling on single and multiple FPGAs

Limitations:

- **Lower simulation throughput** and **longer initial latency**, but still more than sufficient for rapid design-space exploration

CCMT

38

Collapsed Pipeline Single-FPGA Performance/Scalability

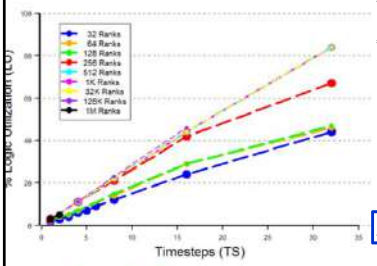

Ranks	TS	Num. of Events	% LU	Latency (cycles)	Hardware MSPS†	Hardware GEPS‡	BE-SST KEPS‡	Hardware Speedup
32	32	43,008	44	7,532	10.5	450	10.14	44×10^6
64	32	92,160	46	10,252	5.23	482	24.91	19×10^6
128	32	190,464	47	10,316	2.62	498	24.42	20×10^6
256	32	393,216	67	13,516	1.31	515	23.92	21×10^6
512	32	811,008	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

Pipelines scale linearly with length of simulation, however blockram eventually becomes limiting resource
- e.g., only fit 2 TS for 1 million ranks on Stratix V
- Motivation to explore partially-collapsed pipeline approach

Increased performance expected with Stratix 10
- Collapsed hardware specifically designed to exploit new Stratix 10 architecture

Now able to reach million+ rank simulations in hardware

Rank limit due to insufficient blockram not logic

Expect greater performance when Stratix 10 becomes available

Throughput > 10Mx higher than BE-SST

CCMT

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

39

Collapsed Pipeline Multi-FPGA Performance/Scalability

Pipelines scale linearly with the length of simulation, but as a single, unidirectional pipe
- Partitioned easily/predictably across any number of connected FPGAs

Given a desired # of simulated ranks, timesteps, granularity (events per timestep), we can predict performance & # of FPGAs

Main point: FPGAs achieve similar scale as BE-SST, but orders of magnitude faster
- 1M ranks, 300+ of sims/second

Ranks	TS	FPGAs	Num. of Events	Lat to First Out (cycles)*	Mega Sims /second*	Giga Events /second*
1K	32	1	1,646,592	21,196	0.33	539
1K	64	2	3,293,184	41,439	0.33	1,077
1K	128	4	6,586,368	81,925	0.33	2,155
32K	32	1	55,443,456	241,868	$1.02E^{-2}$	567
32K	64	2	110,886,912	451,039	$1.02E^{-2}$	1,134
32K	128	4	221,773,824	869,381	$1.02E^{-2}$	2,267
128K	32	2	223,346,688	536,863	$2.56E^{-3}$	571
128K	64	4	446,693,376	942,725	$2.56E^{-3}$	1,142
128K	128	8	893,386,752	1,754,449	$2.56E^{-3}$	2,283
1M	**32**	**16**	**1,799,356,416**	**2,641,321**	**$3.19E^{-4}$**	**575**
1M	**64**	**32**	**3,598,712,832**	**4,234,137**	**$3.19E^{-4}$**	**1,150**
1M	**128**	**64**	**7,197,425,664**	**7,414,540**	**$3.19E^{-4}$**	**2,299**

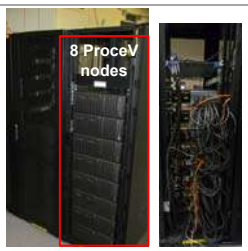

Novo-G#
- 64 GiDEL ProceV (Stratix V)
- 4x4x2 3D torus or 5D hypercube
- 6 Rx-Tx links per FPGA
- Measured 32 Gbps per link
- 150 ns latency across links
- Require 64bits/cycle between FPGAs
- 335MHz, 21.4 Gbps, 51 additional lat
- 450MHz, 28.8 Gbps, 68 additional lat
- 500 MHz before link BW saturation

C. Pascoe, S. Chenna, G. Stitt, H. Lam, "A FPGA-Pipelined Approach for Accelerated Discrete-Event Simulation of HPC Systems", Heterogeneous High-performance Reconfigurable Computing (H2RC), Denver, CO, Nov. 13, 2017.

CCMT

*Predicted collapsed performance of CMT-Bone-BE with varied MPI ranks and simulation timesteps on multiple Stratix V S5GSMD8K1F40C2 @ 335MHz

40

FPGA-Acceleration Summary

- Introduced two approaches for accelerating BE simulations:
  - Fully expanded pipeline: 7-9 orders of magnitude faster than BE-SST, but doesn't scale well
  - Collapsed pipeline: 6 orders of magnitude faster than BE-SST, and scales well across multiple FPGAs
- Demonstrated simulations of 1M ranks
  - Larger than what BE-SST can currently handle
- Integrated symbolic regression into FPGA accelerator performance models
  - Determining attractive tradeoffs between accuracy and scalability is important for FPGA

CCMT | 41

Outline

- Review & overview:
  - Behavioral Emulation (BE)

Research thrusts & achievements

- BE-SST si
- BE metho
  - *Multi-f*
  - *Symbo*
  - *Therm*
- Accelerat
  - Speedu
  - Scalabi

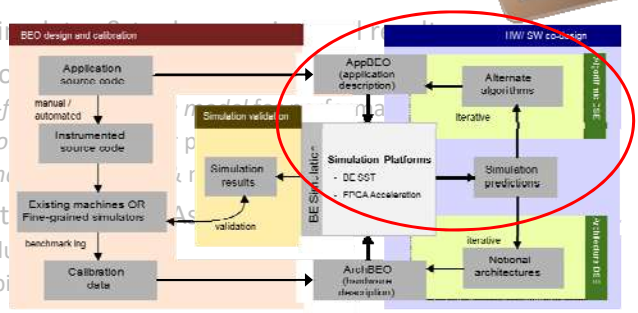

- End-to-end case study - CMT-nek DSE using FPGAs
  - Design space exploration of CMT-nek
- Summary, conclusions, & future work

CCMT | 42

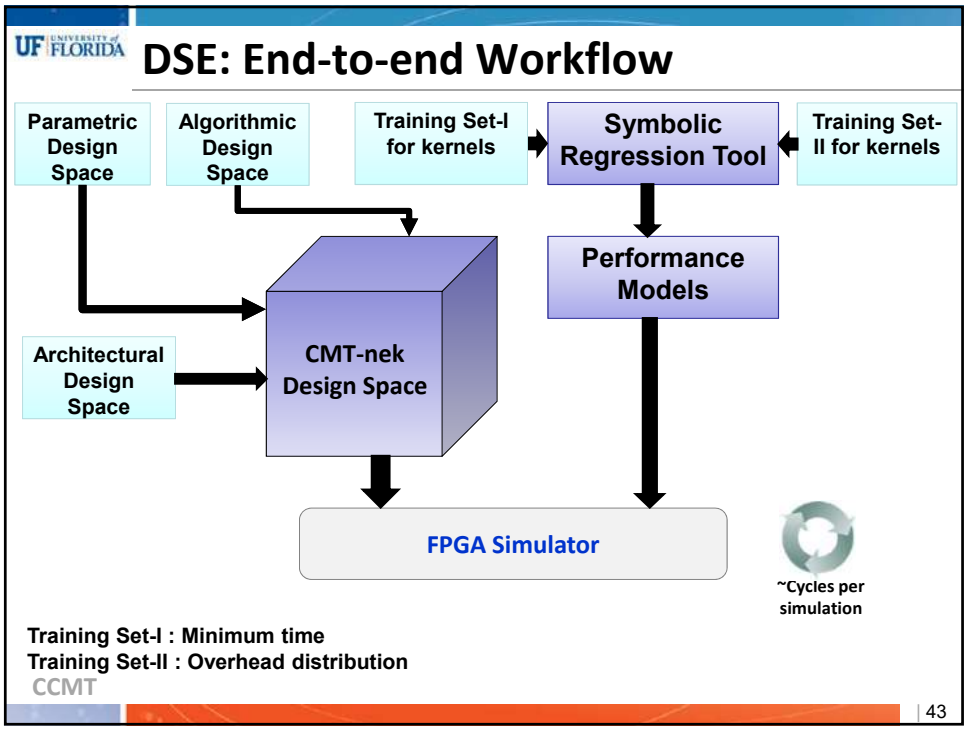

Design Space Exploration

- Goal: Use FPGA to explore enormous design space
- Application: CMT-nek particle solver kernel
- Performance parameters:

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

CCMT

42

Design Space Exploration: Results

Ranks	No of events	FPGA Avg. Simulation time (microseconds)	BE-SST* Avg. Simulation time (seconds)	FPGA speedup
1024	24,576	3.53	8.234	2.33×10^6
2048	49,152	7.06	6.5	9.21×10^5
4096	98,304	14.12	17.3	1.22×10^6
8192	196,608	28.24	31	1.10×10^6
16k	393,216	56.48	120	2.12×10^6
32k	786,432	112.96	145	1.28×10^6
64k	1,572,864	225.92	405	1.79×10^6
128K	3,145,728	451.84	531	1.17×10^6
256K	6,291,456	903.68	1569	1.74×10^6
512K	12,582,912	1807.36	7416	4.10×10^6
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**

CCMT

*ran on HiPerGator - 64 cores

45

Design Space Exploration: Results

- **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)

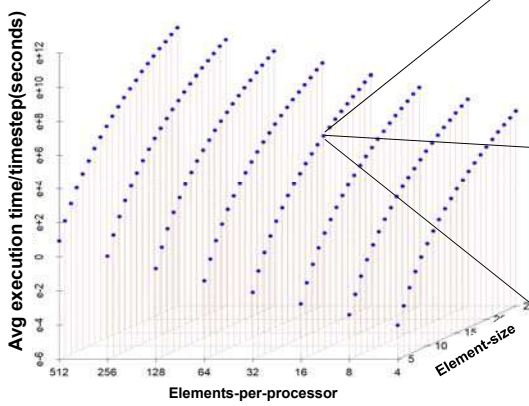

BE-FPGA Simulation

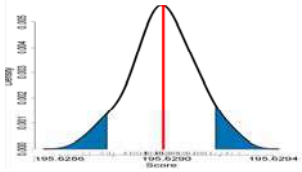

BE-SST Simulation (32k cores)

100 runs &
100 simulations

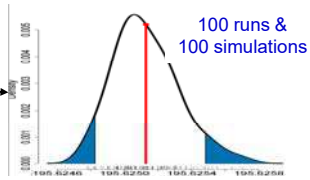

CMT-nek Particle-Solver Execution (4k cores)

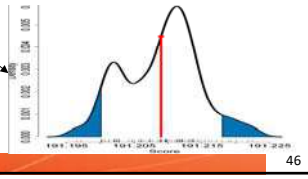

CCMT *ran on HiPerGator - 64 cores **100 Monte-Carlo simulations/candidate

46

Design Space Exploration: Results

BE-FPGA Simulation predictions on Vulcan

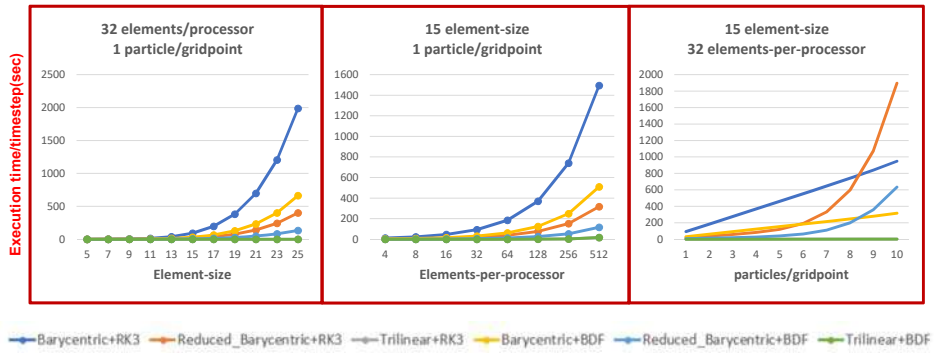

- Algorithms vary
  - ***non-linearly*** w.r.t ***element-size(N)***
  - ***linearly*** w.r.t ***particles/gridpoint(α)*** and ***elements-per-processor(nelt)***
- BDF time-integration provides an average **3x** speedup over baseline* implementation
- Reduced barycentric interpolation provides an average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

CCMT

*baseline = rk3 time integration + barycentric interpolation

47

Key Achievements (Year 4)

- Large-scale experiments using BE-SST
  - Validation (100k+ MPI ranks) on Titan & Vulcan {< 9% error}
  - Predictive simulations approaching a million MPI ranks {Vulcan+ & Titan+}
  - Notional/near-future architecture - SUMMIT performance prediction
- Innovations in BE methods
  - UQ | Symbolic regression | Thermal modeling
- FPGA acceleration of BE simulation
  - Validation of dataflow, pipelined approaches (FEP, CP)
  - Scalability via multiple FPGAs
- End-to-end case study: CMT-nek DSE using FPGA
  - Algorithmic and parametric design space exploration
  - Achieved 10^6 times speedup over BE-SST

CCMT

48

Outline

- Review & Overview:
  - Behavioral Emulation (BE)
- Research thr
  - BE-SST si
  - BE metho
    - *Multi-*
    - *Symbo*
    - *Therm*
  - CMT-nek
  - Acceleration using FPGAs
    - Of BE: speedup vs BE-SST, scale up on multiple FPGAs
    - Of CMT-nek kernels: approach & preliminary results
- Summary, conclusions, & future work

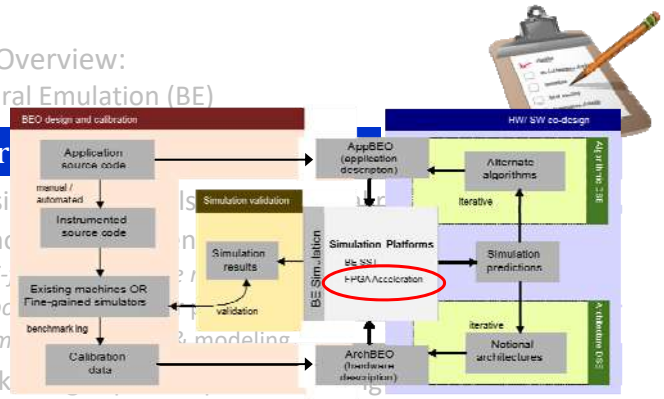

CCMT

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
  - 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

CCMT

Do you have any questions?

CCMT

Appendix

Dynamic Routing

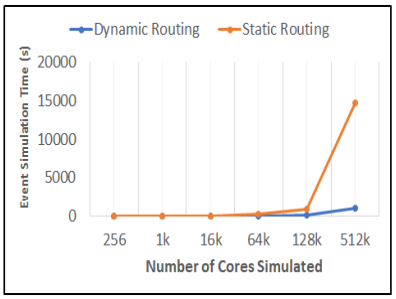

Fig. 1 Static vs. Dynamic Routing

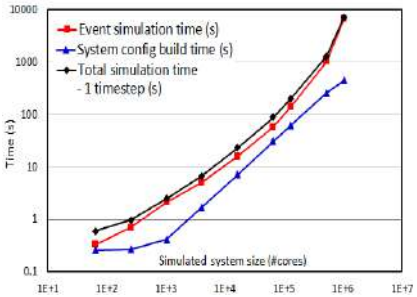

Fig. 2 BE-SST parallelized on 64 cores on HiPerGator @UF

- **Replaced static routing method with dynamic routing**
  - Improve BE-SST scaling: large static routing table not required for increasing number of cores (Fig. 1)
  - Build-time benefits: Config build time not a bottleneck for large simulation (Fig. 2)

CCMT

53

Parallel Performance of BE-SST

- Simulations of Vulcan System with 5-D Torus Topology*
  - 4K, 32K, 128K core
- Improved parallel performance with increase in problem size
  - For e.g., 128k core simulation
    - Without parallelization (1 rank): ~1 day;
    - With 256-rank parallelization: 17 min.

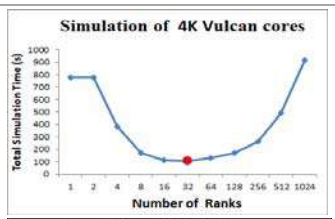

CCMT

* Run using BE-SST installed on Titan

54

BE-SST performance on fat-tree architecture

Comparison Interconnects between 5-D Torus and non-blocking fat tree architecture in BE-SST

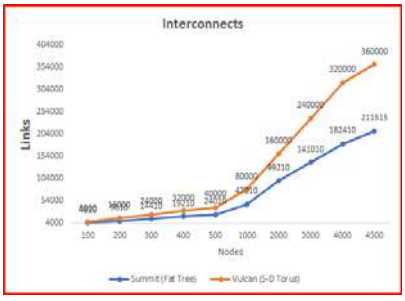

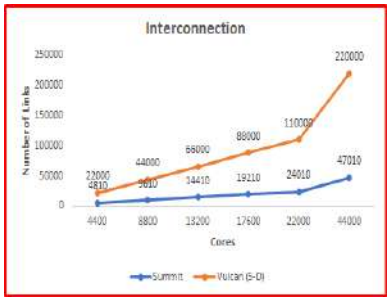

- In simulating equal # of nodes in BE-SST, 5-D torus configuration uses **~1.7 times** the #of interconnects in fat-tree
- In simulating equal # of cores in BE-SST, 5-D torus configuration uses **~4.7 times** the #of interconnects in fat-tree

CCMT

55

BE-SST Performance on Fat-tree Arch

Comparison GCM* between 5-D torus and non-blocking fat tree archs in BE-SST (Build-time)

* **GCM:** Graph Construction Memory: initial memory allocation before simulation

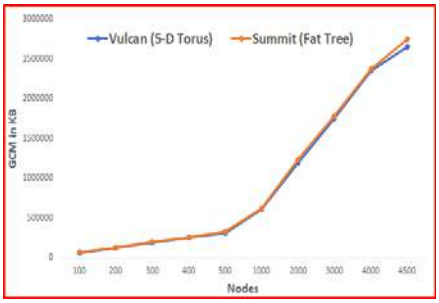

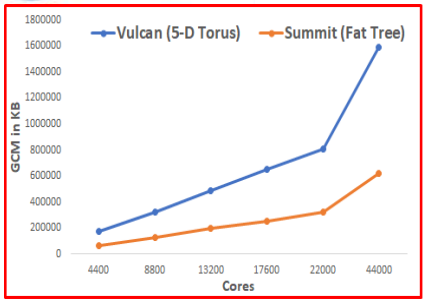

- In simulating equal # of ***nodes*** in BE-SST, GCM* for fat-tree configuration is approx. the ***same*** as the GCM* of 5-D torus
- In simulating equal # of ***cores*** in BE-SST, fat-tree configuration consumes only ***~40%*** of GCM*Fetagu15
- consumed by 5-D torus configuration

CCMT

56

BE-SST performance on fat-tree architecture

Comparison GC time between 5-D Torus and non-blocking fat tree architecture in BE-SST

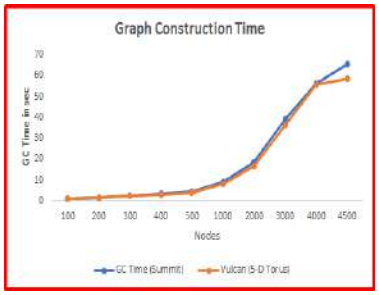

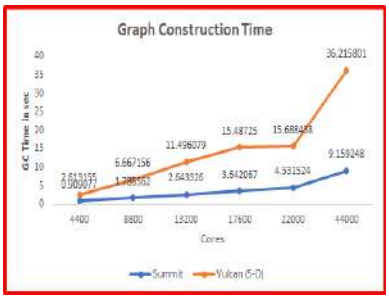

- In simulating equal # of nodes in BE-SST, GCT for 5-D torus configuration is **approx. same** as the GCM of fat-tree topology
- In simulating equal # of cores in BE-SST, 5-D torus configuration consumes **~3.7 times** the GCM of interconnects in fat-tree topology

CCMT | 57

Multi-fidelity Prediction of CMT-nek

w/ UB

Experimental setup

- Design of experiment (DOE)
  - Element size (ES) = 5,9,13,17,21
  - Elements per processor (EPP) = 8,32,64,128,256
  - Number of processors (NP) = 16,256,2048,16384,131072
- 125 total data points
  - ∗ BE simulation: all 125 runs
  - ○ CMT-nek: 22 runs
  - □ CMT-bone: 67 runs

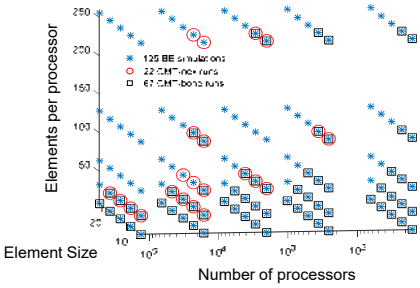

- Multi-fidelity surrogate model
  - Fitting CMT-nek using corrected fitting of BE simulation
  - For large problems, low-fidelity BE simulation is computationally cheaper than high-fidelity CMT-nek or CMT-bone

CCMT | 58

Multi-fidelity Prediction of CMT-nek

w/ UB

Initial RMSE comparison:

- BE-simulation vs CMT-nek at all experimental points
- Overall root mean squared error (RMSE) 74%

Corrected BE simulation against CMT-nek

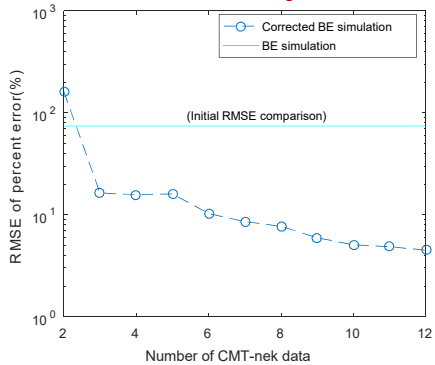

Accuracy of corrected BE sim.

- At 10 left-out CMT-nek test points
- Overall RMSE < 10% with 7 or more CMT-nek data
- Reducing computation budget
  - Supplementing high-fid. (CMT-nek) with low-fid. (BE simulation) runs

Y. Zhang, A. Neelakantan, N. Kumar, C. Park, R. Haftka, N. H. Kim, and H. Lam, "Multi-fidelity Surrogate Modeling for Application/Architecture Co-design", published to Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS17), Denver, CO, Nov. 13, 2017.

CCMT

59

Extrapolation Towards Large-scale Runs

w/ UB

- Extrapolation necessary for exascale prediction
  - But challenging due to limited runs
- Develop MFS using:
  - 7 (out of 22) runs of CMT-nek (high-fidelity), at 256 and 2k processors
  - All 125 BE simulations (low-fidelity)
- Extrapolate using MFS towards 16k ● (11 runs) & 128k ● (4 runs) processors

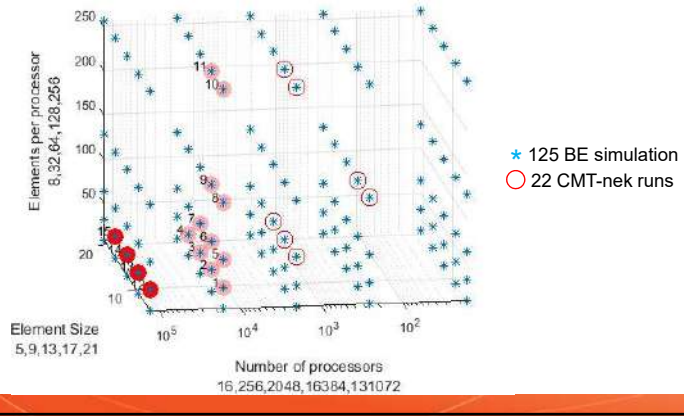

CCMT

60

Extrapolation: Validation

w/ UB

- Extrapolate using MFS towards 16k (11 runs) and 128k (4 runs) processors
  - Linear fit to CMT-nek: poor result
    High-order fitting leads to overfitting due to limited samples (7 CMT-nek runs with 3 variables)
  - MFS: relative errors less than 11%

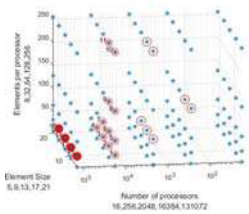

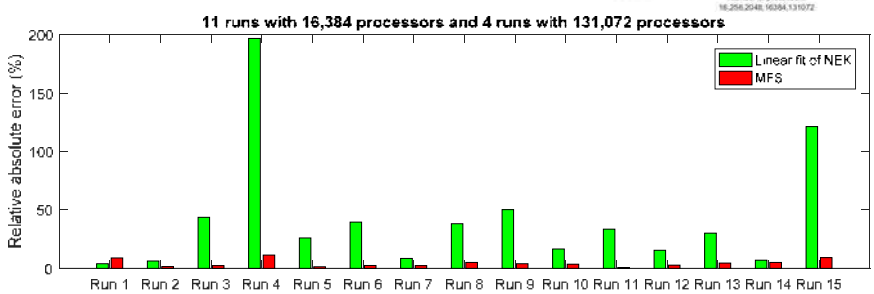

CCMT

61

Extrapolation: Prediction of CMT-nek

w/ UB

- MFS prediction towards one million processors with UQ
  - 64 elements per processors (9 CMT-nek runs)
  - MFS predictions ◯)
- Monotonic trend with small curvature in logarithmic coordinate

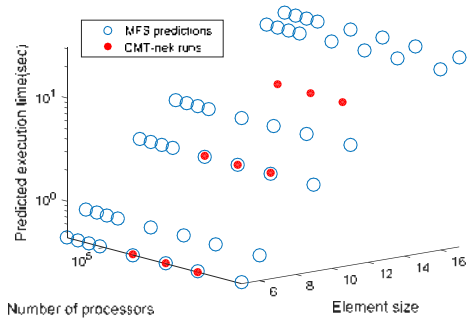

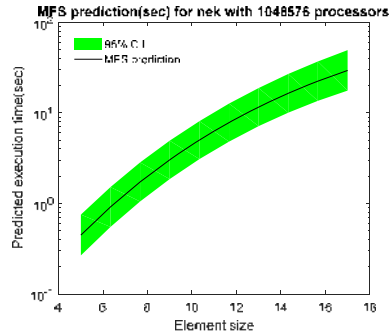

CCMT

62

CCMT

CMT-nek: CS Update

Sanjay Ranka
Computer and Information Science and Engineering
University of Florida
sanjayranka@gmail.com
352 514 4213

UF UNIVERSITY of FLORIDA

Members


Keke Zhai


Mohamed Gadou


Adeesha Malavi


Tania Banerjee


David Zwick


Jason Hackl


Sai Chenna


Sanjay Ranka

CCMT 2

Exascale Simulation: CS Challenges

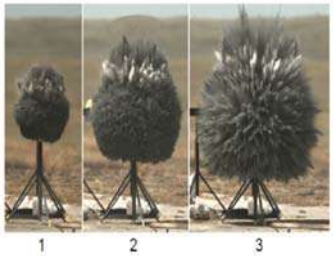

Dynamic Load Balancing

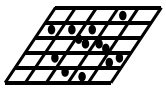

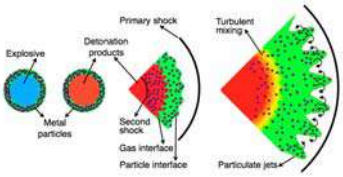

Energy and Performance Tradeoffs

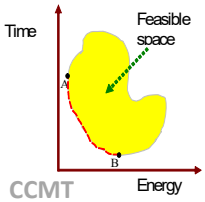

Future Architectures and Hybrid Processors

Parallelization on a million+ processors

CCMT

3

Accomplishments (past year)

- Dynamic load balancing of CMT-nek for moving particles
- Results for CMT-bone with DVFS on heterogeneous platform
- Support for CMT-bone on KNL (architectures with multi level memories)
- Support for CMT-nek on GPU (in progress)

CCMT

4

Talk Overview

- CMT-bone/CMT-nek
- Dynamic Load Balancing
  - Colocation of particles with spectral elements
  - Mapping/Remapping
- Hybrid Multicore Processor
  - Traditional Core (Autotuning)
  - GPU (Manual Tuning)
  - Hybrid Core Mapping
  - KNL (Multi-level Memories)
  - Power/Energy Modeling
  - Dynamic Voltage Scaling
- Conclusions

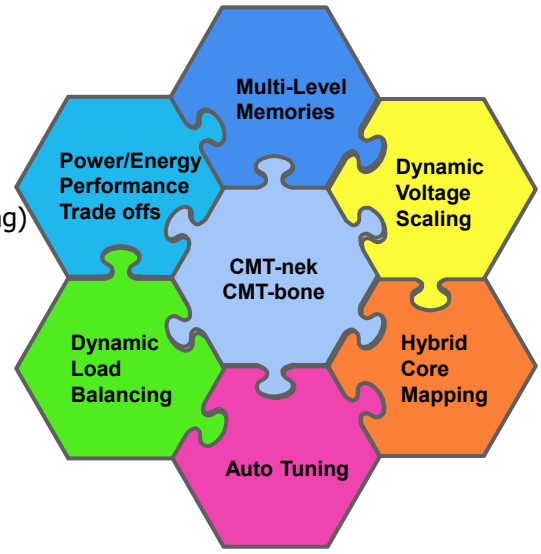

CCMT

5

CMT-bone and CMT-nek workflow

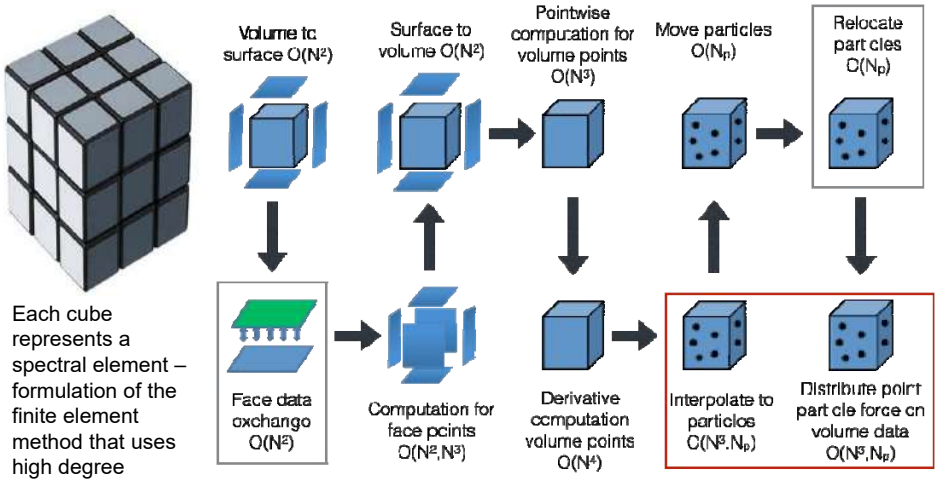

Each cube represents a spectral element – formulation of the finite element method that uses high degree piecewise polynomials as basis functions

CMT-bone -- A Proxy Application for Compressible Multiphase Turbulent Flows, Jason Hackl, Mrugesh Sringarpure, Tania Banerjee, Tanzima Islam, S. Balachandar,Thomas Jackson, Sanjay Ranka, HiPC 2016, December 2016.

CCMT

6

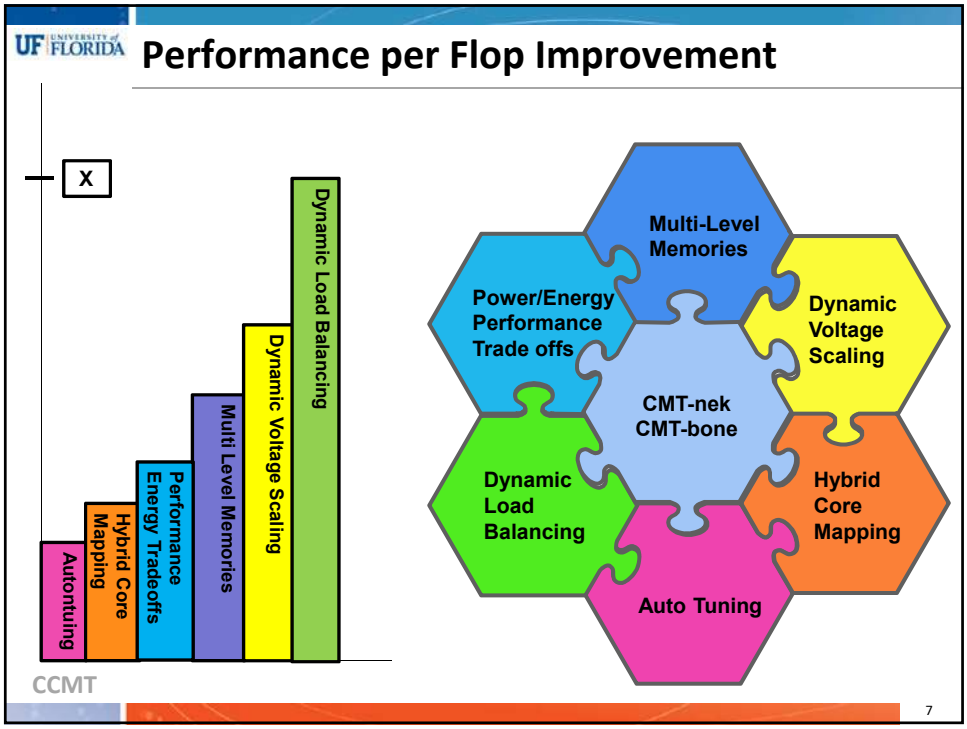
UF UNIVERSITY of FLORIDA
Performance per Flop Improvement
X
Autotuning
Hybrid Core Mapping
Performance Energy Tradeoffs
Multi Level Memories
Dynamic Voltage Scaling
Dynamic Load Balancing
Multi-Level Memories
Power/Energy Performance Trade offs
Dynamic Voltage Scaling
CMT-nek CMT-bone
Dynamic Load Balancing
Hybrid Core Mapping
Auto Tuning
CCMT
7

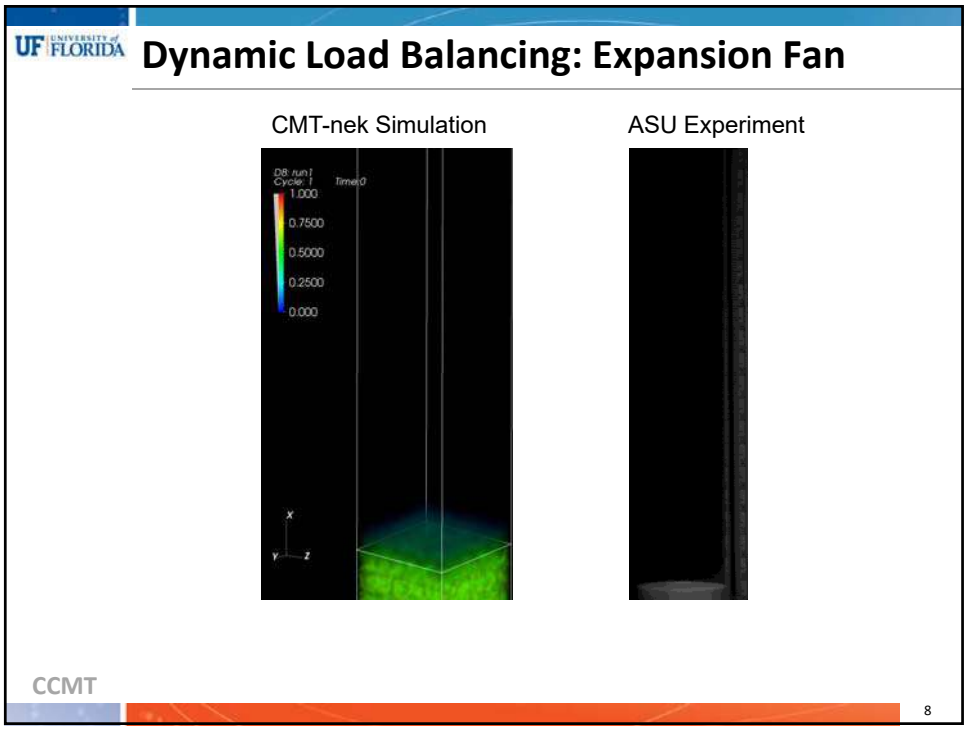
UF UNIVERSITY of FLORIDA
Dynamic Load Balancing: Expansion Fan
CMT-nek Simulation
ASU Experiment
1.000
0.7500
0.5000
0.2500
0.000
CCMT
8

Demonstration Problem

CMT-nek Simulation

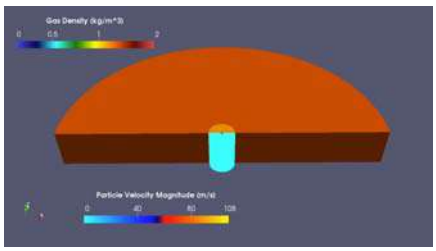

Eglin Blast Pad Experiment

CCMT

9

Overview of Dynamic Load Balancing

Step 1: Domain decomposition

Happens during initialization only

Step 2: Elements to processor mapping

Happens during initialization and on every remapping

Step 3: Decide when to trigger a remap

Rebalance after every k time steps (user set up)

Rebalance automatically after certain time steps (adaptive load balancing)

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

CCMT

10

Spectral Bisection

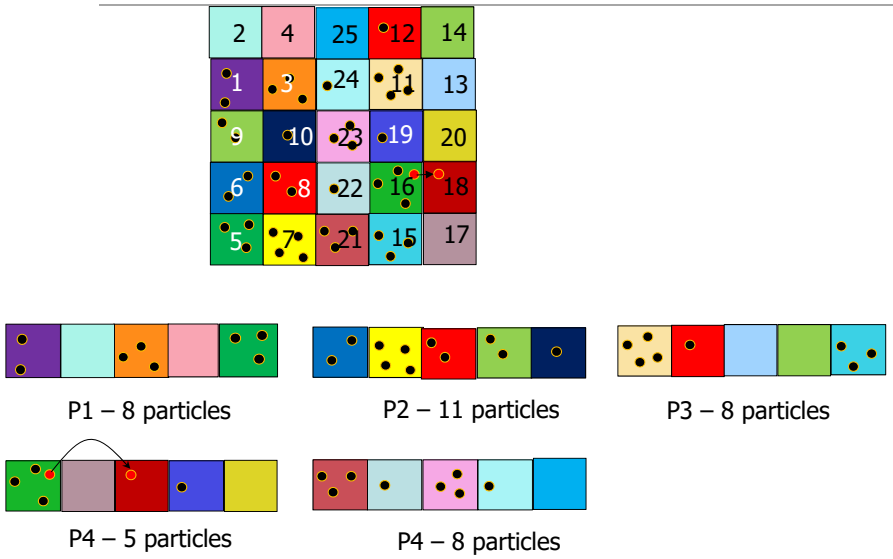

CCMT

11

Elements to Processor Mapping

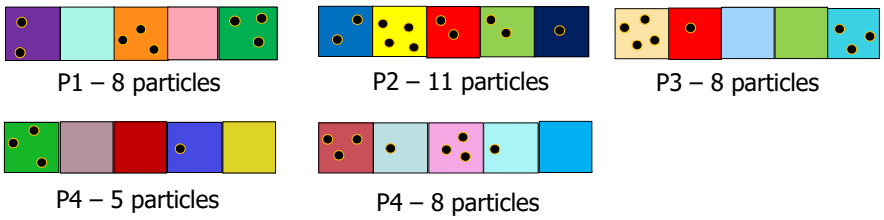

(a) Default elements to processor mapping, agnostic of particles

CCMT

12

Elements to Processor Mapping

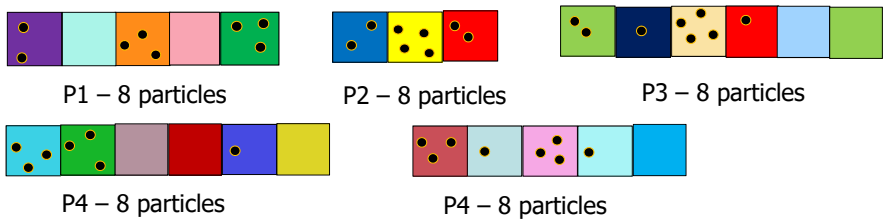

(b) Elements to processor mapping, after taking into account both elements and particles

CCMT

13

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

14

The Centralized Algorithm

Each processor sends element load array to P0

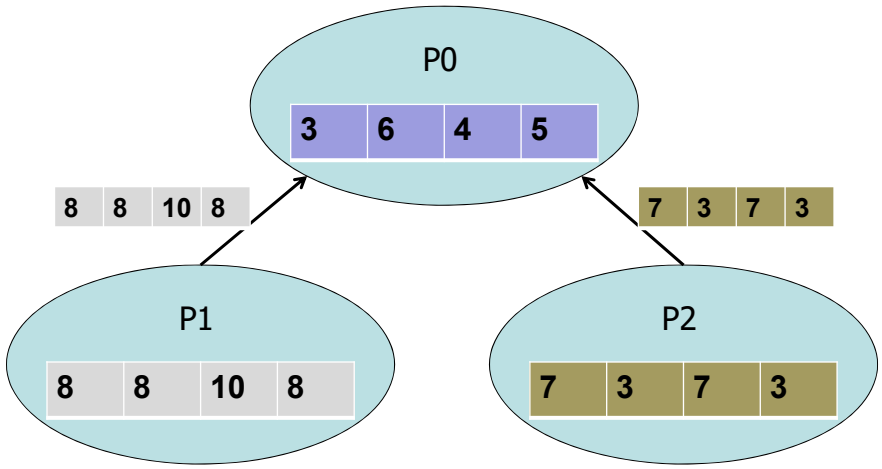

CCMT 15

The Distributed Algorithm

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

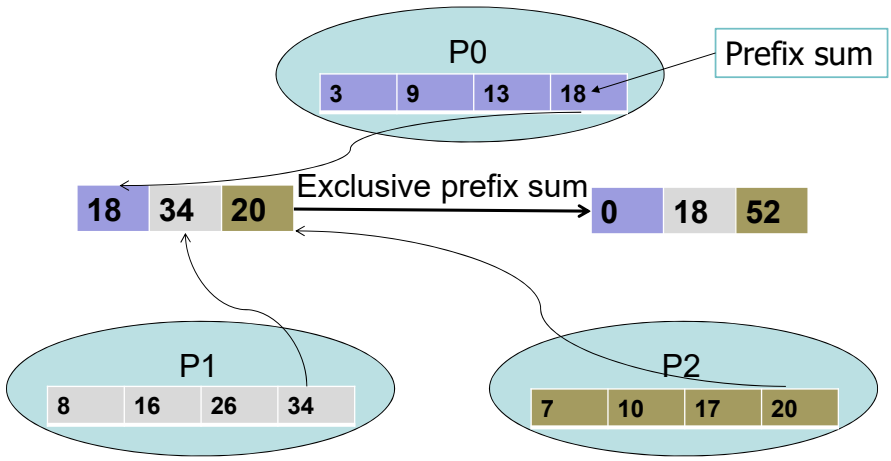

CCMT 16

The Hybrid Algorithm

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

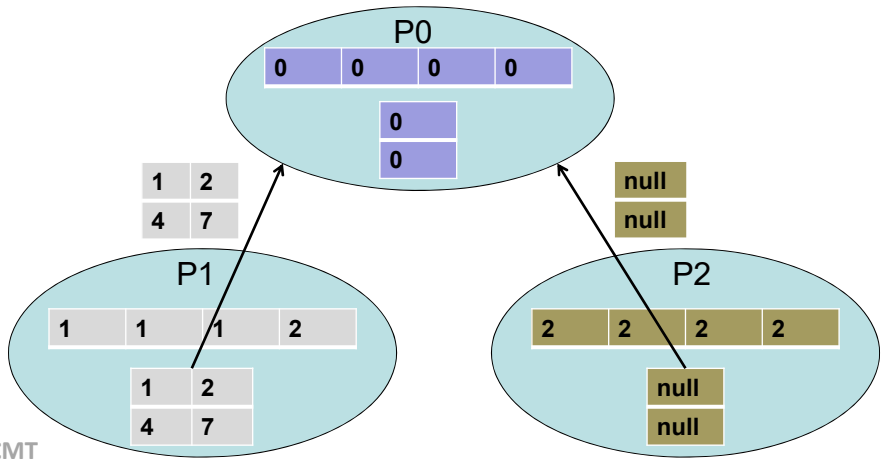

CCMT

17

Rebalancing Time: Total Overhead Time (Vulcan)

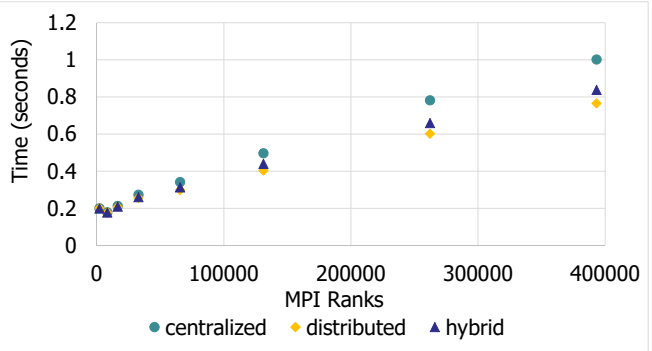

CMT-nek

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
  - 0.77 for the distributed,
  - 0.84 for the hybrid algorithm

CCMT

18

Expansion testcase: CMT-nek on Vulcan

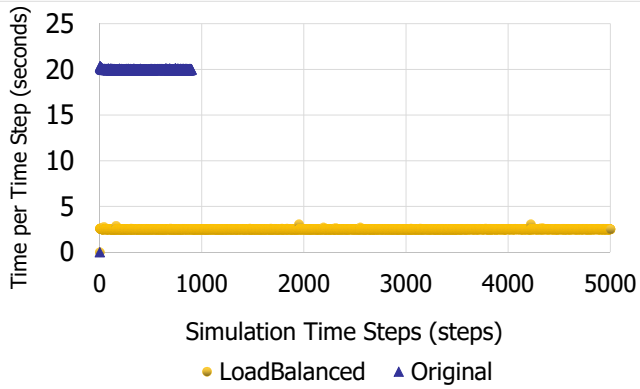

CMT-nek

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

- Time per time step:
  - 20.00 s for the original code
  - 2.52 s for the load balanced code
- 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.

CCMT

19

Expansion testcase: Adaptive Load Balancing

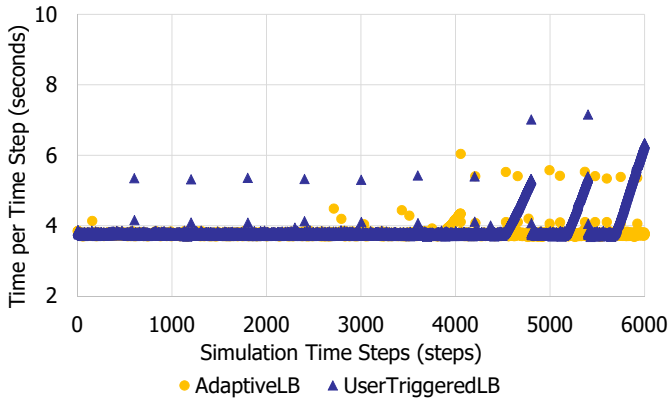

CMT-nek

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%.

CCMT

20

Performance comparison between CMT-nek and Nek5000

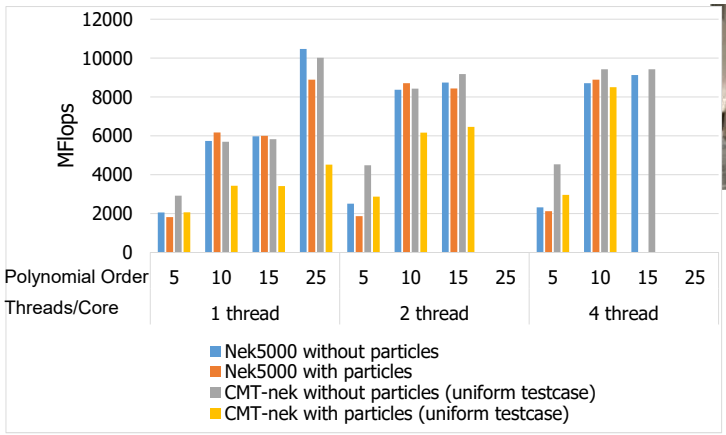

CMT-nek

16 MPI ranks
1 nodes
64 elements
16384 particles
Uniform test

- Performance of CMT-nek without particles is about the same as Nek5000 without particles
- Performance of CMT-nek with particles is about half of the performance of Nek5000 with particles

CCMT

21

Performance of CMT-nek with particles using multiple-thread

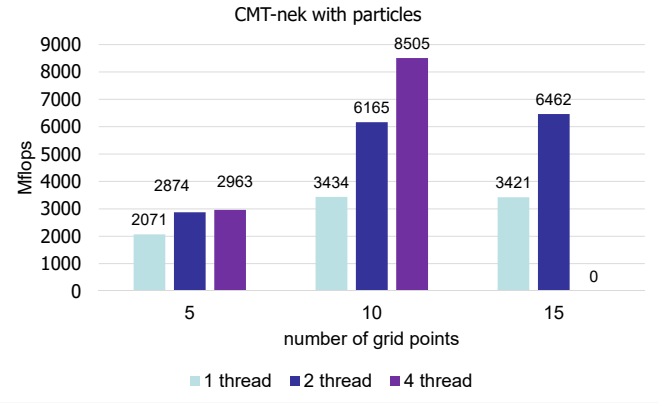

CMT-nek

16 MPI ranks
1 nodes
64 elements
16384 particles
Uniform test

- Performance using 2 threads is about 2 time performance using 1 thread

CCMT

22

Expansion testcase: Power consumption on Mira (Using MonEQ)

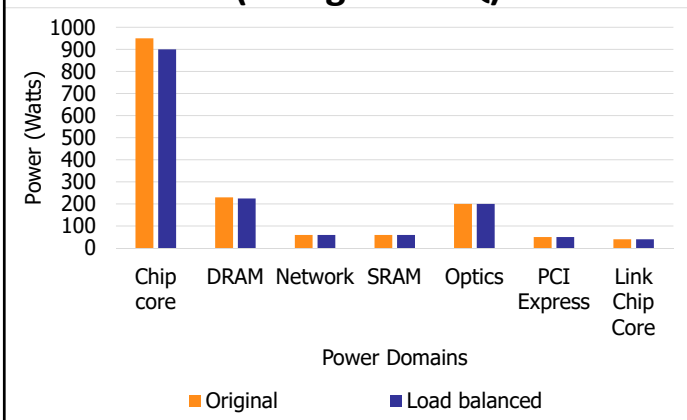

CMT-nek

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

CCMT

23

Rebalancing Time: Total Overhead (Quartz)

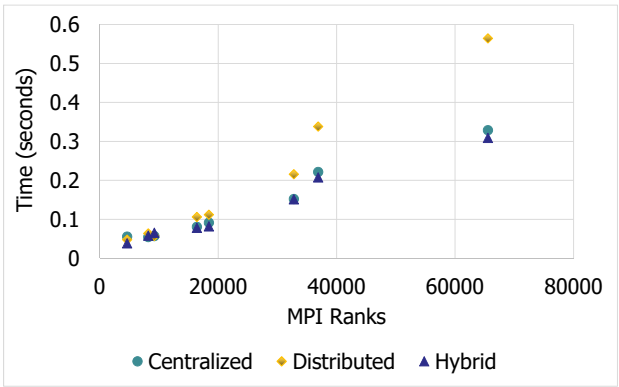

CMT-nek

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
  - 3.35 for the distributed
  - 1.82 for the hybrid algorithm
- This shows that the overhead for load balancing is very small.

CCMT

24

Expansion testcase: CMT-nek on Quartz

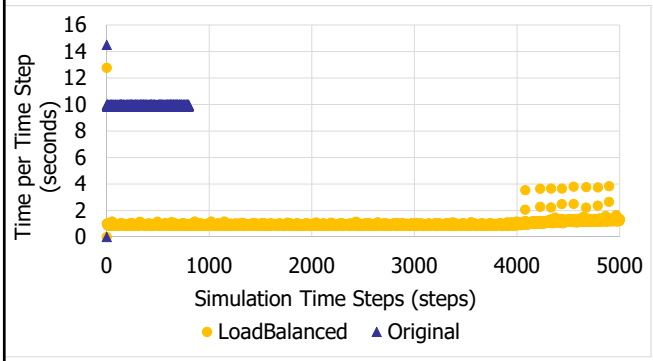

CMT-nek

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

- 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.

CCMT

25

Expansion testcase: Power consumption on Quartz (using Libmsr)

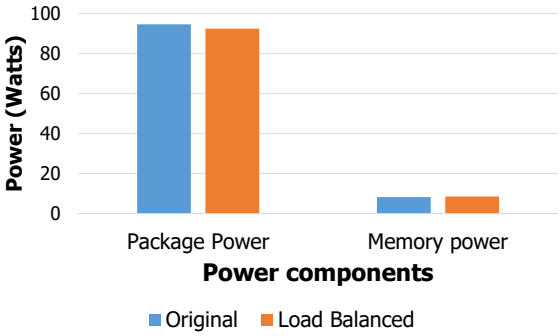

CMT-nek

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

- Power consumption is comparable.

CCMT

26

Nek5000 User Meeting

- To be held at Tampa, on April 17 and 18.
- Load balancing will be on interest since Nek5000 is being enhanced with particles.
- Potentially, there will be a wider user base for load balancing.

CCMT

27

Hybrid Multicore Architectures

Multi-objective Load decomposition

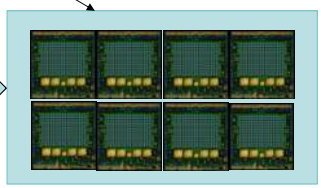

CPU Attribute	Description
Frequency	2.4GHz
Sockets	2
Cores/Socket	12

GPU Attribute	Description
Model	K40m
Frequency	745MHz
SMs	15
Cores/SM	192

Multi-objective Optimization on Hybrid Systems using DVFS, Mohamed Gadou, Sankeerth Mogili, Tania Banerjee, Sanjay Ranka, submitted to ICCP 2018

CCMT

28

CPU: Power-Performance Tradeoffs

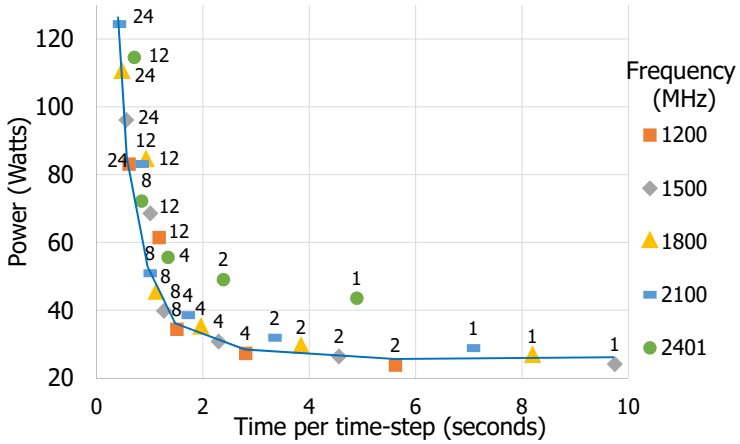

Single node
CMT-bone
8192 elements
128000 particles
7x7x7 grid size
Strong scaling

- Observations:
  - Using 24 cores causes very high power consumption
  - Using 8 and 4 cores is power as well as performance efficient

CCMT

29

CPU: Energy-Performance Tradeoffs

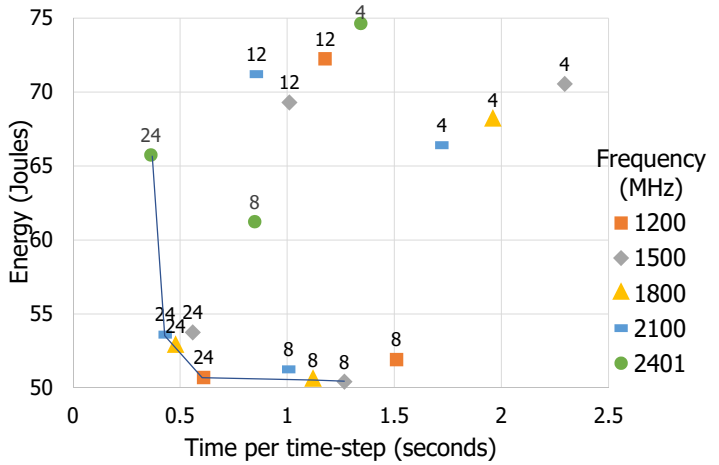

Single node
CMT-bone
8192 elements
128000 particles
7x7x7 grid size
Strong scaling

- Observations:
  - Using 24 cores with lower frequencies results in an energy and performance efficient configuration

CCMT

30

GPU: Power-Performance Tradeoffs

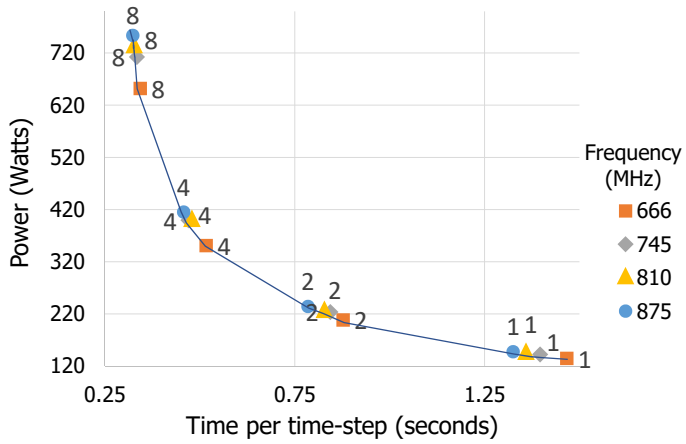

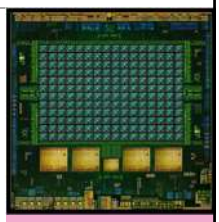

Single node
CMT-bone
8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Observations:
  - Using 8 GPUs causes very high power consumption
  - Using 4 and 2 GPUs is power as well as performance efficient

CCMT

31

GPU: Energy-Performance Tradeoffs

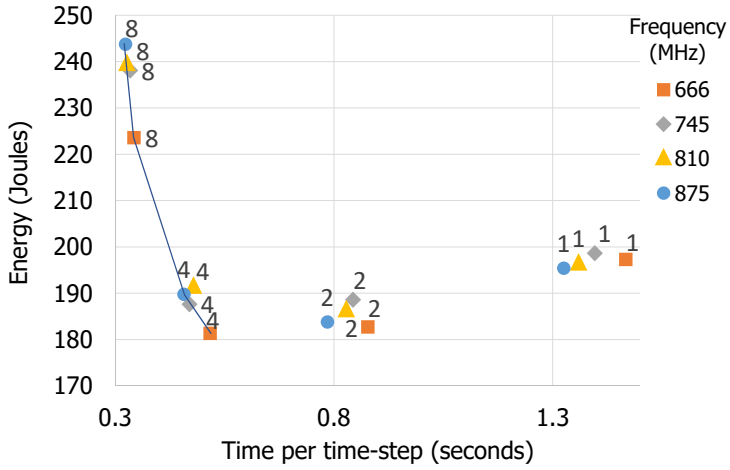

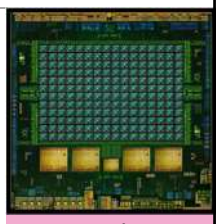

Single node
CMT-bone
8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Observations:
  - Using 4 GPUs with lower frequencies results in an energy and performance efficient configuration

CCMT

32

Theoretical formulation

Minimize $E(n, m, f, g, S)$

Subject to $T(n, m, f, g, S) <= T_1 - O(n, m, f, g, S)$

$n <= N \rightarrow 24$

$m <= M \rightarrow 8$

$f \in F \rightarrow 12$

$g \in G \rightarrow 7$

16128 S

X → n CPU cores, freq f

S-X → m GPU cores, freq g

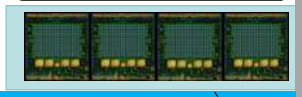

Minimize $E_{CPU}(n, f, X) + E_{GPU}(m, g, S-X)$

Subject to $T_{CPU}(n, f, X) <= T_1 - C_o(m+n, S)$

$T_{GPU}(m, g, S-X) <= T_1 - C_o'(m, S-X)$

$n <= N \rightarrow 24$

$m <= M \rightarrow 8$

$f \in F \rightarrow 12$

$g \in G \rightarrow 7$

$T_{CPU}(n, f, X) \times P_{CPU}(n, f, X)$ — 344S — $T_{GPU}(m, g, S-X) \times P_{GPU}(m, g, S-X)$

$T_{CPU}(n, X)/f \times P_{CPU}(n, f)$ — 344+32S — $T_{GPU}(m, S-X)/g \times P_{GPU}(m, g)$

E : Energy consumed
n : Number of CPUs
m : Number of GPUs
f : CPU Frequency
g : GPU Frequency
S : Problem size
T : Computation Time
T_1 : Constrained time
O: Communication overhead
N : Total available CPUs
M : Total available GPUs
F : Set of available CPU frequencies
G : Set of available GPU frequencies

33

CPU Experiments

Parameters	Values
x, problem size in terms of number of elements (57)	1024, 1152, 1280, 1408, 1536, 1664, 1792, 1920, 2048, 2176, 2304, 2432, 2560, 2688, 2816, 2944, 3072, 3200, 3328, 3456, 3584, 3712, 3840, 3968, 4096, 4224, 4352, 4480, 4608, 4736, 4864, 4992, 5120, 5248, 5376, 5504, 5632, 5760, 5888, 6016, 6144, 6272, 6400, 6528, 6656, 6784, 6912, 7040, 7168, 7296, 7424, 7552, 7680, 7808, 7936, 8064, 8192
f, frequency in MHz (5)	1200, 1500, 1800, 2100, 2400
n, number of cores (5)	1, 2, 4, 8, 12

CCMT

34

CPU Experiments

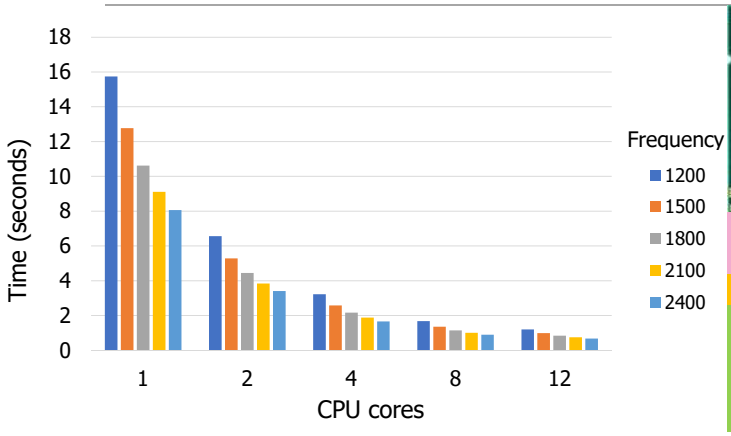

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Time per time step for processing a problem size of 8192 spectral elements and 256,000 particles
- Experiments show time decreases by increasing CPU counts and frequency

CCMT

35

CPU Experiments

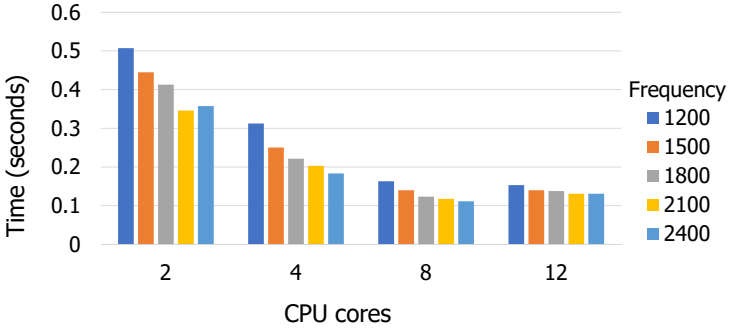

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- CPU communication time for processing 8192 spectral elements and 256,000 particles
- Communication time first decreases and then increases as processors are increased because volume of data communicated per core decreases in the beginning but the total communication volume begins to increase later

CCMT

36

Assumptions

- Performance using lower clock frequencies may be derived from that using the highest clock frequency
  - $T_{CPU}(n, f, x) = k \times \frac{f_{max}}{f} \times TCPU(n, fmax, x)$
- Power consumption does not depend on problem size
  - $P_{CPU}(n, f, x) = PCPU(n, f)$
- n = CPU count, f = CPU frequency, x = problem size
- Example: Power lookup table:

f \ n	1	2	4	8	12
1200	93.24	97.33	101.32	114.46	127.50
1500	96.67	100.78	108.32	122.95	136.79
1800	98.14	102.88	111.53	131.86	148.45
2100	100.69	106.31	117.11	138.72	160.84
2400	102.65	108.12	123.94	145.54	167.08

CCMT

37

Performance Prediction Error

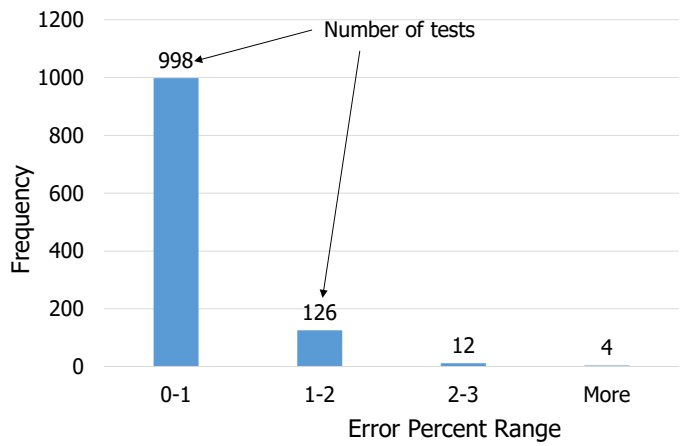

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Prediction error is within 3% for 99% of the tests
- $T_{CPU}(n, f, x) = k \times \frac{f_{max}}{f} \times TCPU(n, fmax, x)$

CCMT

38

Power Prediction Error

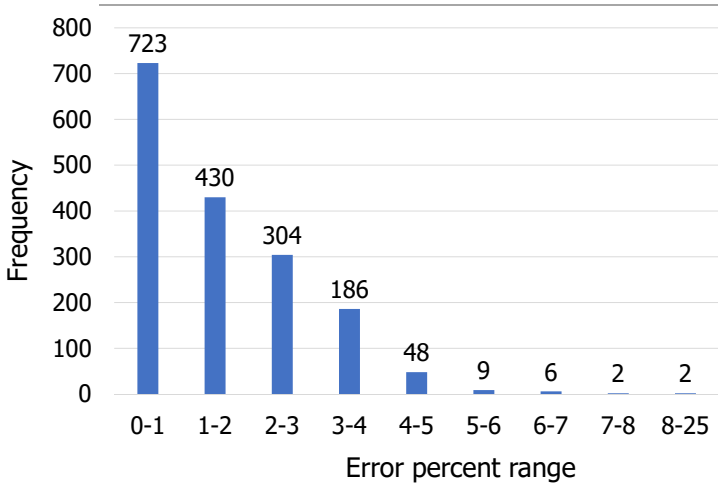

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Prediction error is within 5% for 99% of the tests
- $P_{CPU}(n, f, x) = PCPU(n, f)$

CCMT

39

GPU Experiments

Parameters	Values
y, problem size in terms of number of elements (57)	Same problem sizes as described for CPUs.
g, frequency in MHz (3)	875, 810, 745
m, number of GPUs (4)	1, 2, 4, 8

CCMT

40

GPU Experiments

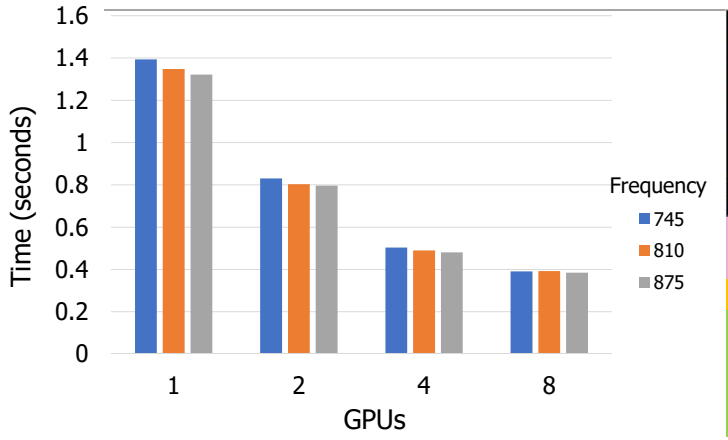

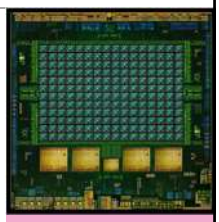

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Time per time step for processing a problem size of 8192 spectral elements and 256,000 particles
- Experiments show time decreases by increasing CPU counts and frequency

CCMT

41

GPU Experiments

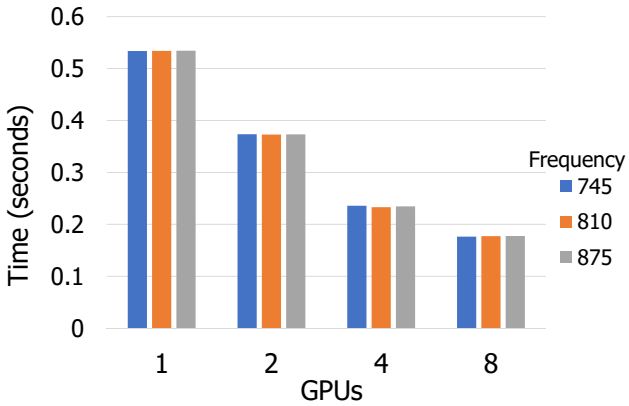

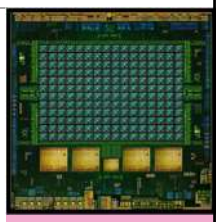

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Total communication time when GPUs are used to solve the problem

CCMT

42

Assumptions

- Performance using lower clock frequencies may be derived from that using the highest clock frequency
  - $T_{GPU}(m,g,y) = k \times \frac{g_{max}}{g} \times TGPU(m,gmax,y)$
- Power consumption does not depend on problem size
  - $PG_{PU}(m,g,y) = PGPU(m,g)$
- m = GPU count, g = GPU frequency, y = problem size
- Example: Power lookup table:

g \ m	1	2	4	8
745	89.79	169.91	314.53	550.91
810	93.51	176.05	324.59	567.11
875	98.80	185.46	341.26	589.41

CCMT

43

Performance Prediction Error

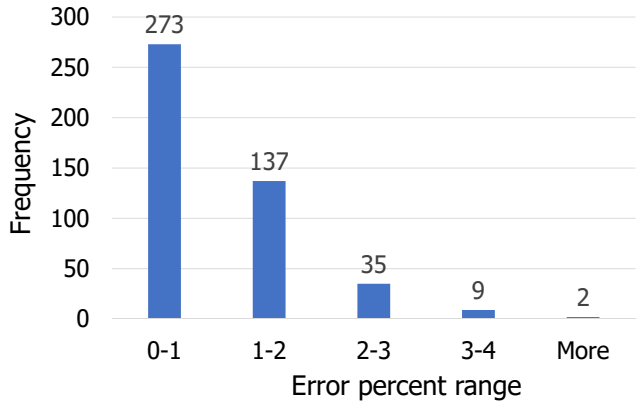

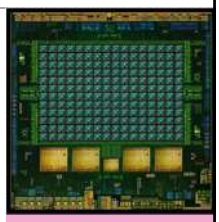

Single node
CMT-bone
8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Prediction error is within 5% for 99% of the tests

CCMT

44

Power Prediction Error

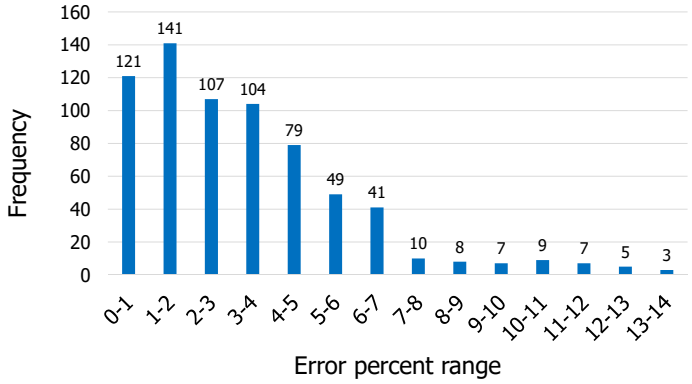

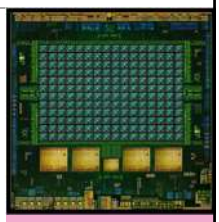

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Prediction error is within 10% for 99% of the tests

CCMT

45

Summary of CPU and GPU experiments

- Assumption that performance at lower clock frequencies may be derived from the performance of the application at the highest clock frequency holds both for CPUs and GPUs
- Assumption that power consumption does not depend on problem size holds for both CPUs and GPUs
- The above assumptions help to reduce the number of required experiments further

CCMT

46

Hybrid Experiments

Parameters	Values
Load (x, y) x = CPU load y = GPU load	(1024, 7168), (1152, 7040), (1280, 6912), (1408, 6784), (1536, 6656), (1664, 6528), (1792, 6400), (1920, 6272), (2048, 6144), (2176, 6016), (2304, 5888), (2432, 5760), (2560, 5632), (2688, 5504), (2816, 5376), (2944, 5248), (3072, 5120), (3200, 4992), (3328, 4864), (3456, 4736), (3584, 4608), (3712, 4480)
f	1200, 1500, 1800, 2100, 2400
g	745, 810, 875
n	2, 4, 8, 12
m	1, 2, 4, 8

CCMT

47

Hybrid Experiments

CPU			GPU			Predicted Energy (J)	Actual Energy (J)	Error %
elements	**count**	**GHz**	**elements**	**count**	**MHz**			
2540	8	2.4	5652	2	875	225.25	232.87	3.27
2621	8	2.4	5571	2	810	226.05	235.01	3.81
2621	8	2.4	5571	2	875	227.13	237.77	4.48
2540	8	2.4	5652	2	810	228.61	227.14	0.65
2293	8	2.4	5899	2	875	229.86	227.72	0.94
1474	4	2.4	6718	2	875	230.35	236.19	2.47

- Error in prediction for some of the best predicted configurations for minimum energy consumption

CCMT

48

Hybrid Experiments

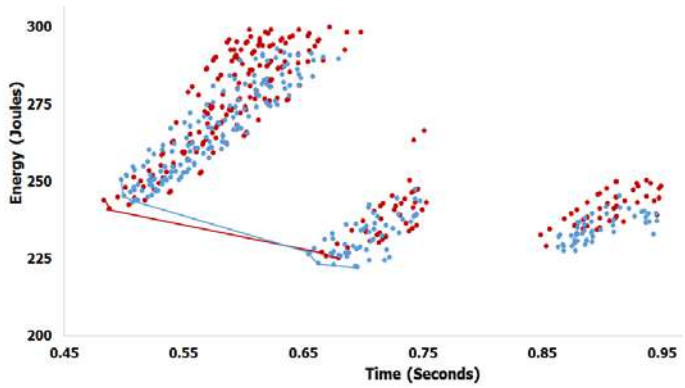

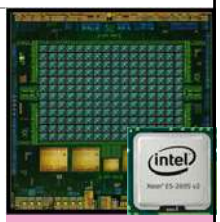

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Energy versus performance pareto optimal curve. Red line shows the predicted curve, while blue line is obtained from actual data

CCMT

49

Pareto Curves for Energy Versus Runtime

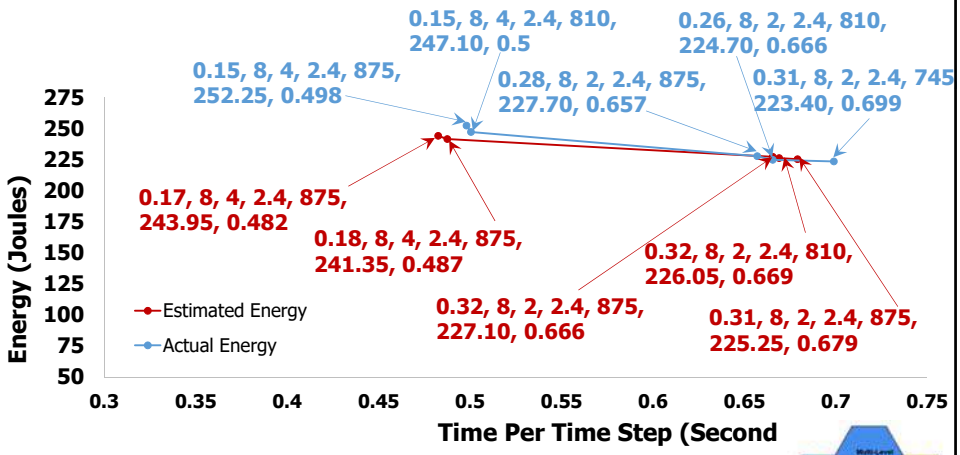

- Data for each Pareto-optimal point contain the following information in order:
  CPU load (%), CPU count, GPU count, CPU frequency (GHz), GPU frequency (MHz), energy (J), time (s)

CCMT

50

Pareto Curves for different input sizes

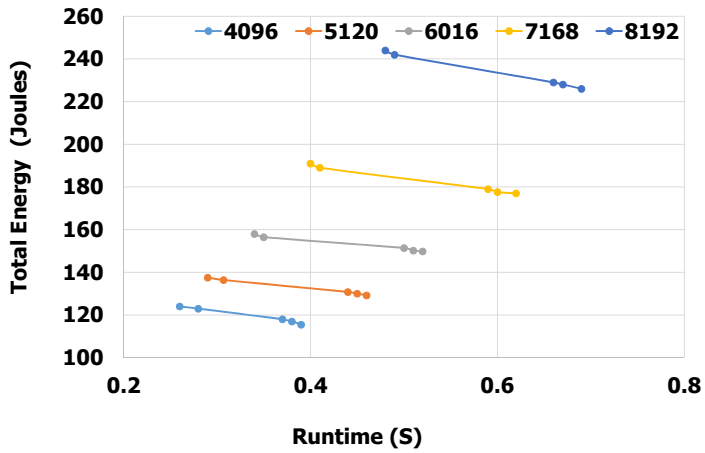

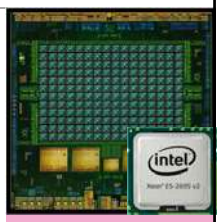

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

CCMT

51

Pareto Curve for Power Vs Runtime

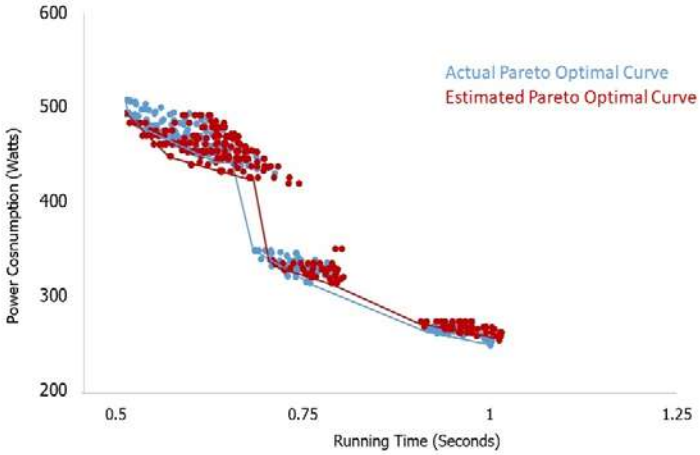

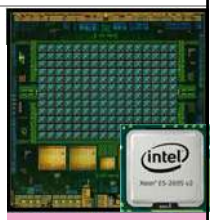

Single node

CMT-bone

8192 elements
256,000 particles
7x7x7 grid size
Strong scaling

- Power versus performance pareto optimal curve. Red line the predicted curve, while blue line is obtained from actua

CCMT

52

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

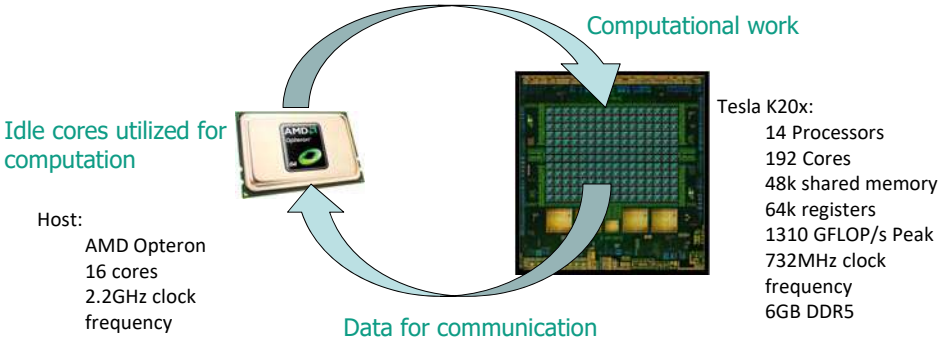

CCMT

53

CMT-nek on GPU

- Based on the implementation of CMT-bone on GPU for fluid, with support for new enhancements (viscous flow)
- 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

CCMT

54

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.
- 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

CCMT

55

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 CPU	4 CPUs	16 CPUs	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

CCMT

56

Multi-Level Memories

KNL Overview

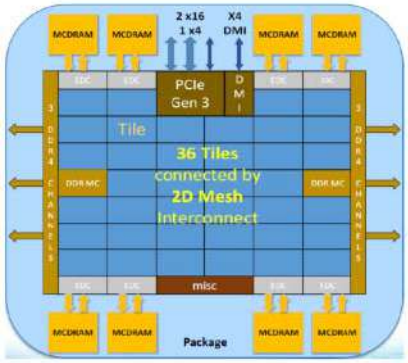

Chip: 36 Tiles interconnected by **2D Mesh**
Tile: 2 Cores + 2 VPU/core + 1 MB L2

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

CCMT

57

CMT-bone on Intel KNL

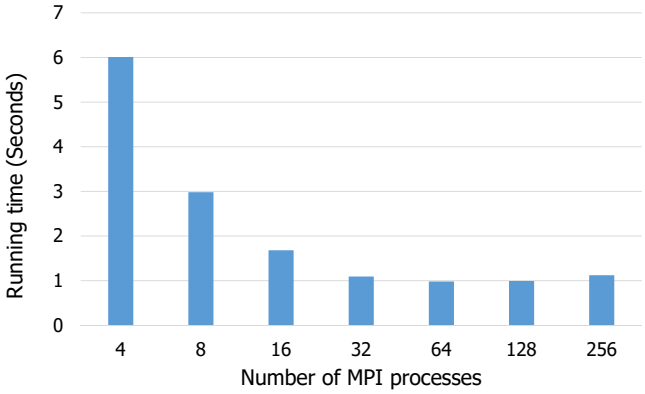

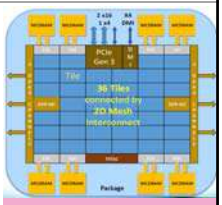

Single node
CMT-bone
7x7x7 grid size
1.8GHz frequency
32 MPI ranks
2 OpenMP threads

- Running time is total time taken to complete 1 time step.
- Running time as Number of processes being varied.

CCMT

58

CMT-bone on Intel KNL (OpenMP)

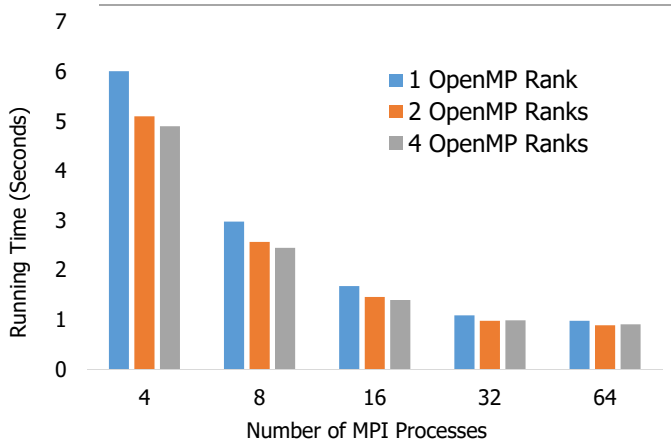

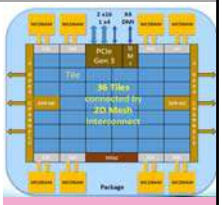

Single node
CMT-bone
7x7x7 grid size
1.8GHz frequency
32 MPI ranks
2 OpenMP threads

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

CCMT

59

CMT-bone on Intel KNL (MCDRAM)

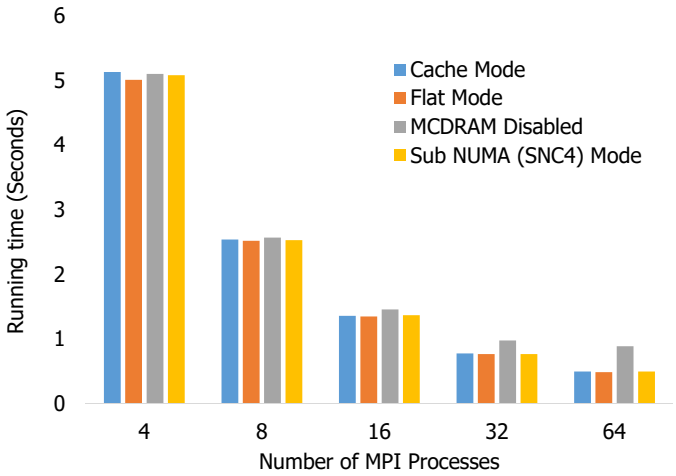

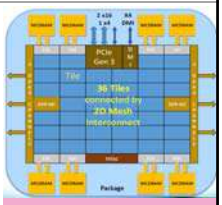

Single node
CMT-bone
7x7x7 grid size
1.8GHz frequency
32 MPI ranks
2 OpenMP threads

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

CCMT

60

CMT-bone on Different Platforms

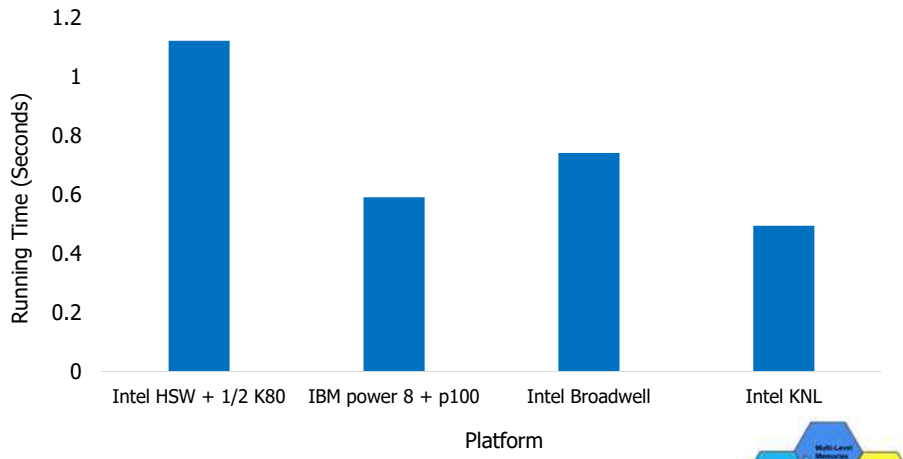

- Only 1 cpu + 1 gpu was used for GPU based platforms
- Intel Broadwell was a 16 core processor
- KNL is faster than Intel Broadwell and gpu accelerated Systems

CCMT

61

Summary of Accomplishments

- Dynamic load balancing of CMT-nek for moving particles
- Results for CMT-bone with DVFS on heterogeneous platform
- Support for CMT-bone on KNL (architectures with multi level memories)
- Support for CMT-nek on GPU (in progress)

CCMT

62

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, submitted to ICPP 2018
10. Mohamed Gadou, Meena Arunachalam, Tania Banerjee, Sanjay Ranka, Multiobjective evaluation and optimization of CMT-bone on Intel Knights Landing, in preparation.

63

Conclusions: Overall Improvement

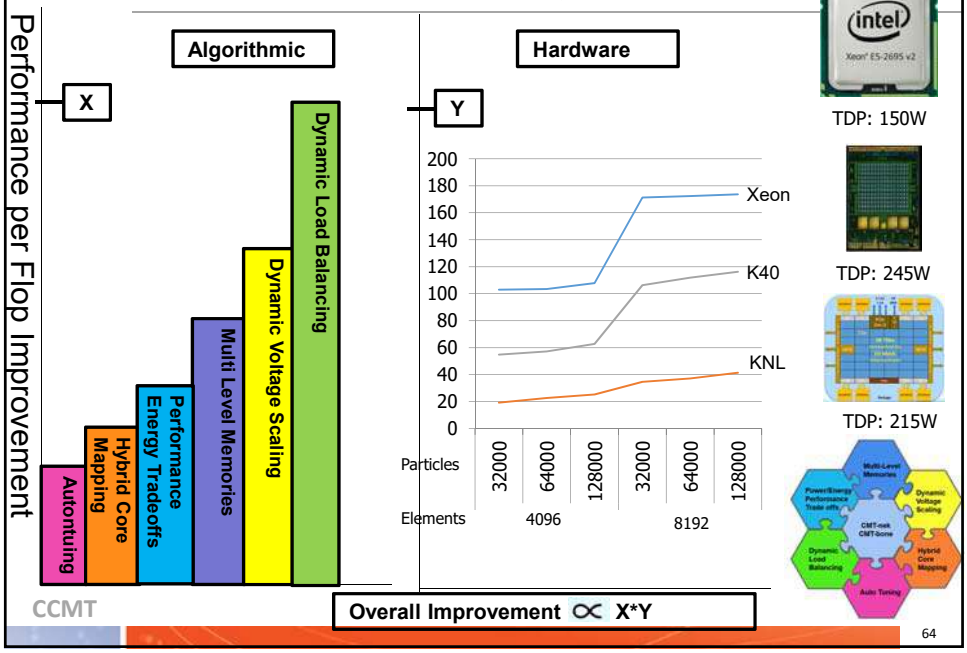

64

CCMT

Do you have any questions?

UF UNIVERSITY of FLORIDA

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

CCMT

66

Hardware-Software Autotuning

Algorithm: *dudr-4loop*

```
do k = 1, Nz
 do j = 1, Ny
  do i = 1, Nx
   do l = 1, Nx
    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, Nz* Ny
 do i = 1, Nx
  do l = 1, Nx
    dudr(l, k) = dudr(l, k) +
                 a(i, l) * u(l, k, ie)
 enddo
enddo
enddo
```

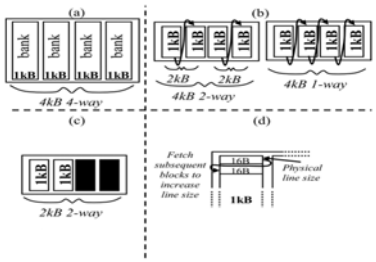

Total number of variants = 20

Algorithm: *dudr-4loop-permuted-and unrolled*

```
do k = 1, Nz
 do i = 1, Nx
  do j = 1, Ny
   do l = 1, Nx,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
```

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

CCMT

67

Genetic Algorithm

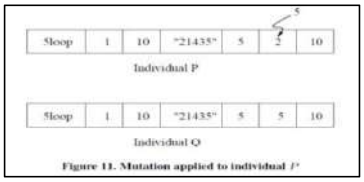

i=1

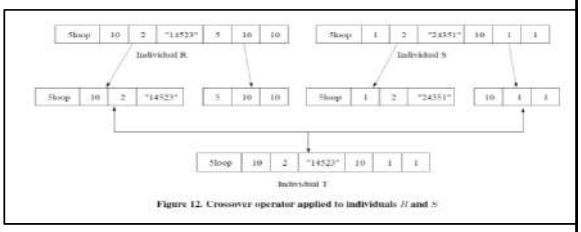

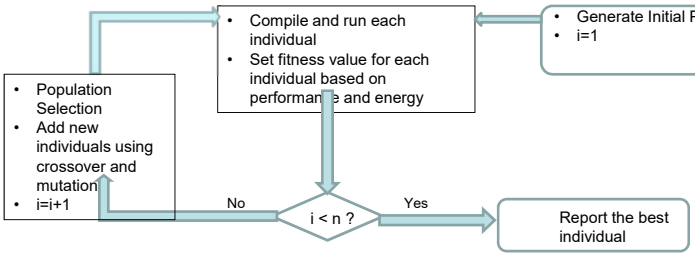

CCMT

68

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)

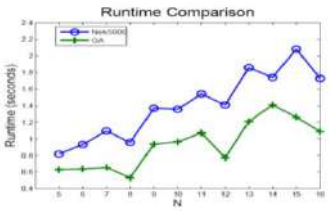

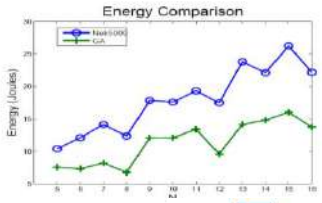

CCMT

Performance (varying hardware parameters)

Best Time (seconds)

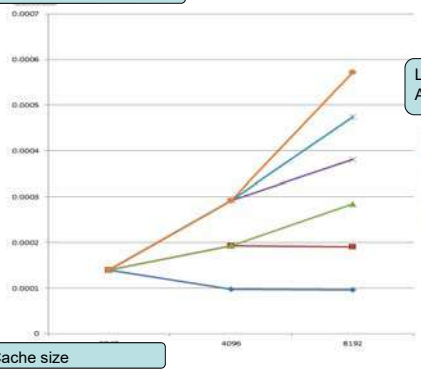

Cache size

Line Size - Associativity

Associativity Tuning

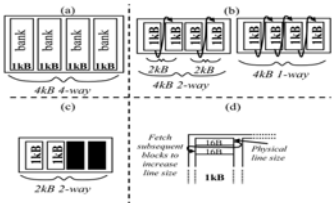

Capacity Tuning

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

70

CCMT

Uncertainty Budget

Validation and Uncertainty Reduction

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

Department of Mechanical & Aerospace Engineering

UF UNIVERSITY of FLORIDA

Uncertainty Budget Team


Raphael Haftka, PhD


Nam-Ho Kim, PhD


Chanyoung Park, PhD


Kyle Hughes


Samaun Nili


Yiming Zhang


Giselle Fernandez


Tanner Johns


Eric Maltz

(New undergrad students)

- PIs: Raphael Haftka and Nam-Ho Kim
- Research Scientist: Chanyoung Park
- Students
  - Kyle Hughes (UQ of Eglin exp.)
  - Samaun Nili (Mesoscale 1D/2D force model potential errors)
  - Yiming Zhang (UQ and multi-fidelity correction of BE emulation)
  - Giselle Fernandez (Instability signal to noise ratio)
  - Tanner Johns (UQ of ASU experiments)
  - Eric Maltz (UQ of Eglin exp.)

CCMT

2

Summary of Past Work

- Year 1: Integrate with physics team
- Year 2: Integrate with exascale team
  - UQ for SNL simulation and experiments
  - Initial UQ for CS and exascale activities
- Year 3: Engage with the demonstration problem
  - Forensic UQ for Eglin and SNL experiments
  - Validation by blind prediction for CS and exascale
- Year 4: Engage with the demonstration problem
  - UQ for demonstration problem
  - Finalizing UQ and validation for shock-tube simulation and experiment

CCMT 3

Outline

- **Unique aspects of UQ at CCMT**
- UQ of demonstration problem
- VVUQ workflow of demonstration problem
- Summary of Sandia Shocktube Campaign
- Forensic UQ of micro- & meso-scale experiments

CCMT 4

Unique Aspects of UQ at CCMT

- Separated treatment of error from uncertainty
  - Uncertainty reduction to reveal model error
  - Error reduction through model update
- Global sensitivity analysis to identify model error potentials
- Forensic UQ
- Measurement processing uncertainty
- Method of multiple lines for extrapolation
- Multi-fidelity surrogate modeling

CCMT 5

Reducing Uncertainty to Expose Errors

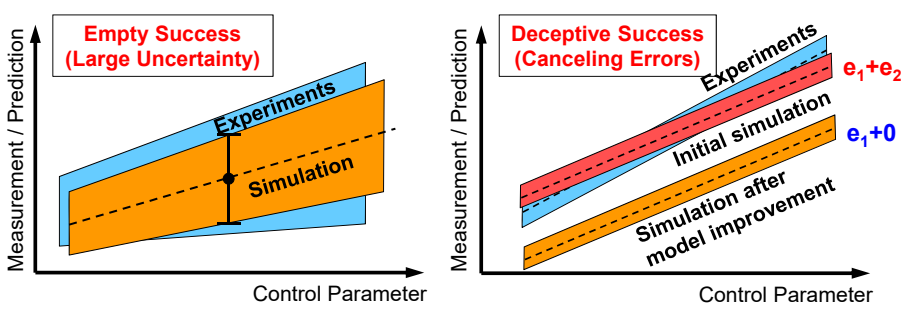

- Reducing uncertainty in model error estimate (UR)
  - Large uncertainty prevents observation of model error
  - Will require substantial uncertainty reduction
  - Hidden model error due to canceling errors
  - Identify models where refinement matters and possibly experimental bias
- Reducing model error (ER)
  - Identifying/ranking error sources

CCMT 6

Forensic UQ

- In the course of digging deeper into experiments and simulations we realized that we perform **forensic uncertainty quantification**
  - Simulation driven UQ: quantifying **uncertainty in measurement** of inputs for simulation
  - Discover anomalies in simulations from UQ of the simulation
  - Identify/help to discover **unrecognized** uncertainties and errors based on UQ
- Forensic UQ tools

Document all the details (crime scene)

Clarify details with experimentalists and physicists (witnesses)

Quantify uncertainties (forensics lab)

Identify unrecognized uncertainties and errors via simulation (culprits)

CCMT

7

Outline

- Unique aspects of UQ at CCMT
- **UQ of demonstration problem**
- VVUQ workflow of demonstration problem
- Summary of Sandia Shocktube Campaign
- Forensic UQ of micro- & meso-scale experiments

CCMT

8

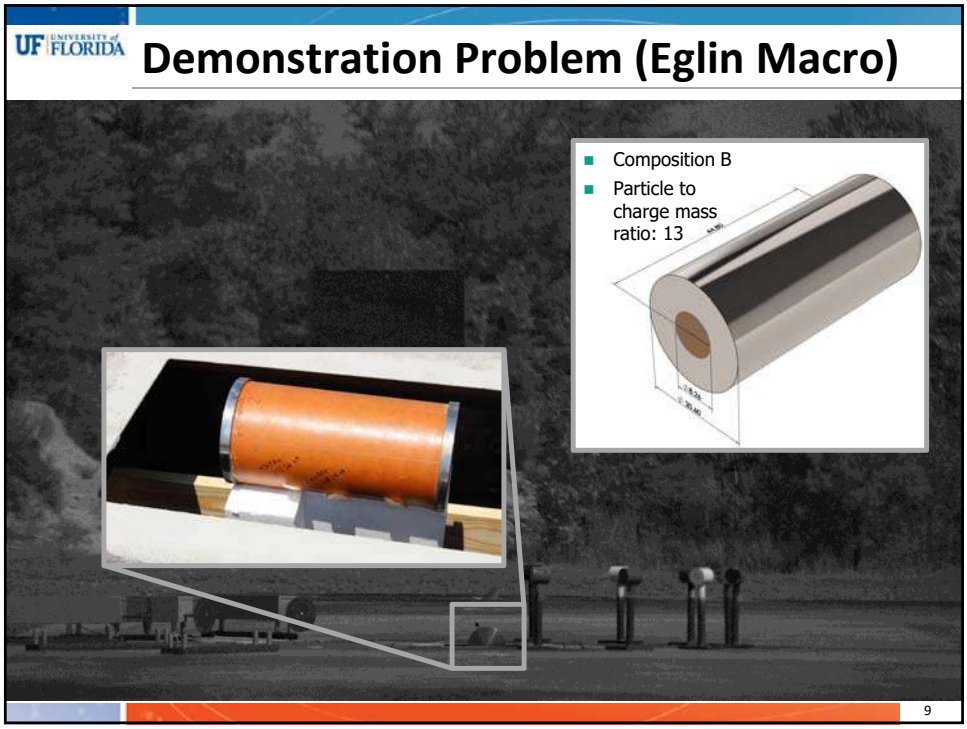
UF UNIVERSITY of FLORIDA
Demonstration Problem (Eglin Macro)
Composition B
Particle to charge mass ratio: 13
9

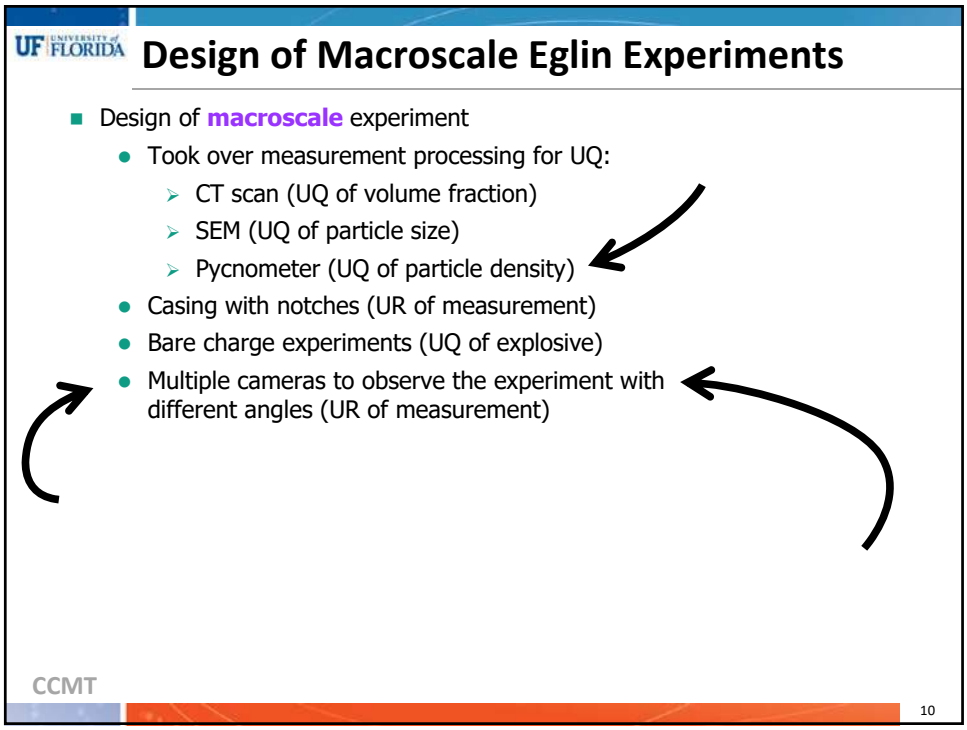
UF UNIVERSITY of FLORIDA
Design of Macroscale Eglin Experiments
Design of macroscale experiment
Took over measurement processing for UQ:
CT scan (UQ of volume fraction)
SEM (UQ of particle size)
Pycnometer (UQ of particle density)
Casing with notches (UR of measurement)
Bare charge experiments (UQ of explosive)
Multiple cameras to observe the experiment with different angles (UR of measurement)
CCMT
10

UQ of Experimental Parameters

Parameter	Quantity	Method
Explosive Length	44.75 ± 0.08 cm	Tape Measure
Explosive Diameter	8.194 ± 0.008 cm	Caliper
Explosive Mass	4100 ± 24 g	Mass Balance
Particle Diameter	TBD	SEM
Particle Density	7.66 ± 0.03 g/cm^3	Pycnometer
Particle Volume Fraction	TBD	CT Scanner
Ambient Pressure	101.8 ± 0.8 kPa	Eglin Weather Station
Ambient Temperature	32 ± 7 °C	Eglin Weather Station
Probe Locations	±1%	Tape Measure

- **Forensic UQ**
  - UB team collected the uncertainty information
  - The information was used for planning the experiment

CCMT

11

UQ Driven Experiment Planning

- Steel particles chosen for validation experiments (75-125 μm)
- To closely match monodisperse and spherical assumptions of the simulation
- Allocated 3 experiments with steel / 2 bare charge / 1 with tungsten
- Particle bed characterization
  - SEM of particles and particle density via pycnometer
  - Steel mock CT scan

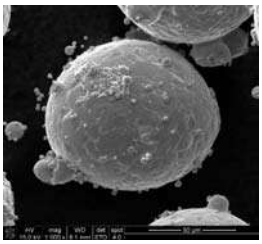

SEM of single steel particle at 1000x zoom

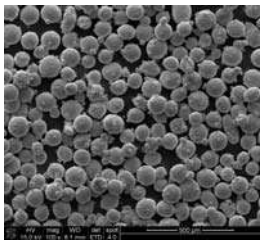

SEM of several steel particles at 100x zoom

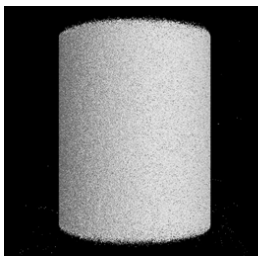

Steel mock CT scan

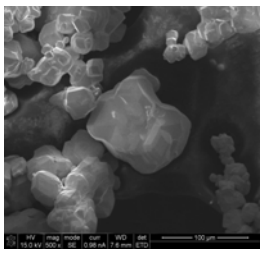

SEM of single tungsten particle, 500x zoom

12

UQ-Driven Casing Design

- Casing fracture can be a possible cause of jet formation
- In order to minimize/control the casing effect (**uncertainty reduction**), two measures were taken:
  - 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)

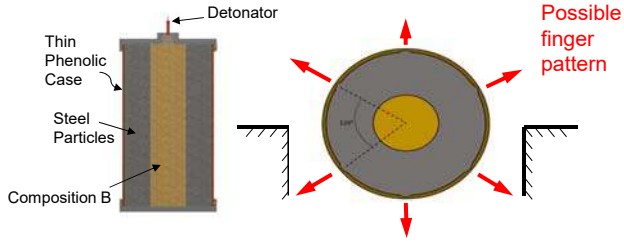

CCMT

13

Negligible Effect of Notched Casing

- 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 were formed for all cases

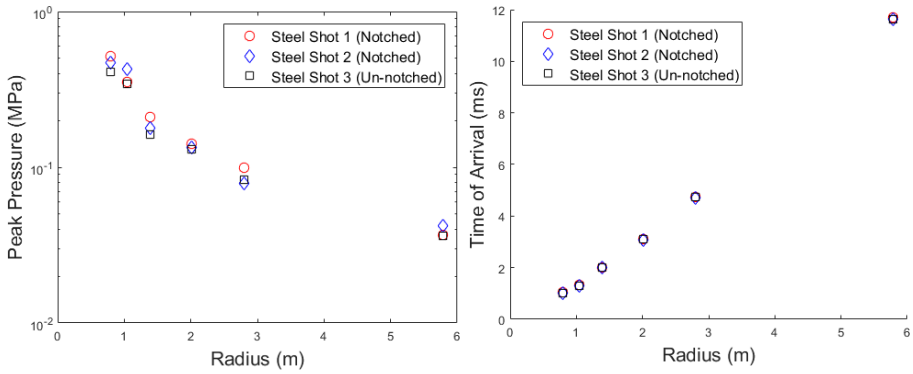

CCMT

14

UQ-Driven Measurement Plan

- Phantom v1212 sampled at 12000 fps (Camera 1/4) and Phantom v711 sampled at 7500 fps
- 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
- Tests performed on AFRL Blastpad (**Kyle Hughes talk**)

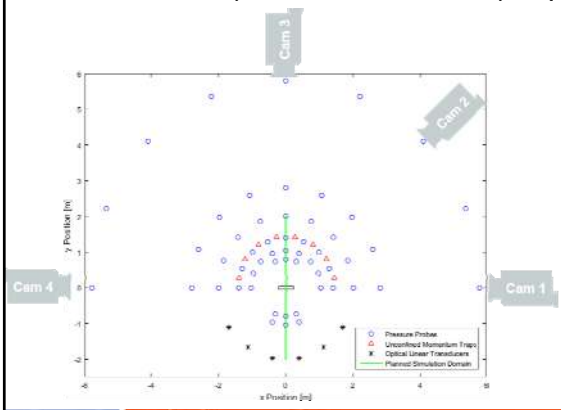

15

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

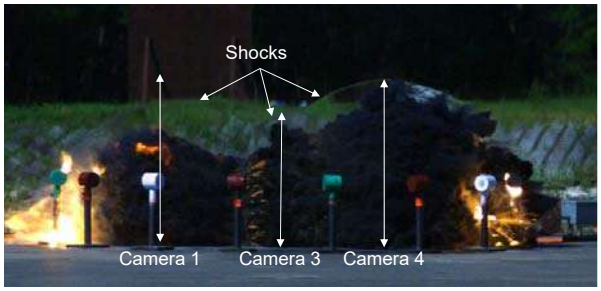

Camera 3 at 1.067 ms after detonation

16

U in CMT-nek Hero 1 Pred. due to T6

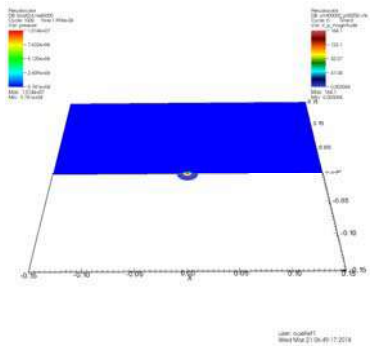

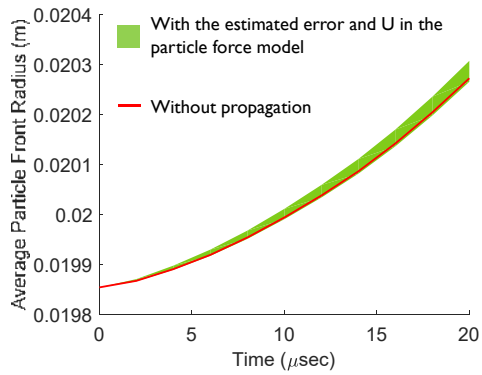

- Error in the **particle force model** (**T6**) was estimated from the Sandia campaign
- Propagating the error and uncertainty in the particle force model through **CMT-nek Hero 1** provides the error and uncertainty in the macroscale prediction
- Errors from other campaigns will be propagated when they are completed

CCMT

17

Outline

- Unique aspects of UQ at CCMT
- UQ of demonstration problem
- **VVUQ workflow of demonstration problem**
- Summary of Sandia Shocktube Campaign
- Forensic UQ of micro- & meso-scale experiments

CCMT

18

V&V and UQ Workflow

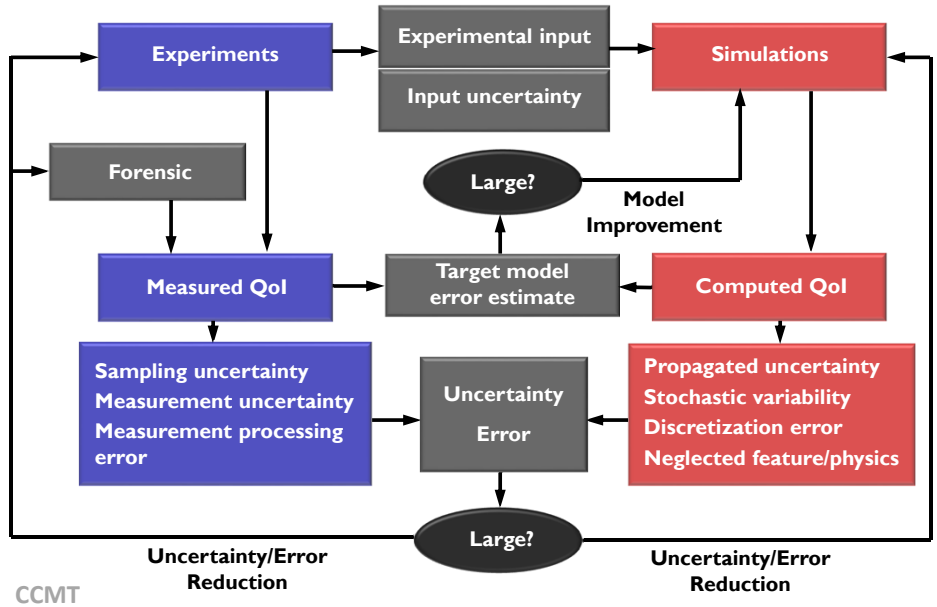

CCMT

19

Discretization Error Estimation

- **Goal:** Estimate the discretization error of multiphase problems
- **Difficulty:** Grid size can be smaller than the particle size
- **Solution:** Use the finite size of the particles (Samaun Nili)

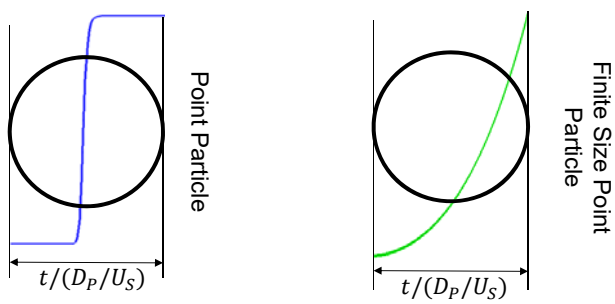

- Point particle shows an abrupt change in force
- Finite size particle shows a gradual increase in force

20

Convergence of Quasi-Static Force

- Can't afford to have the converged grid size, but can estimate discretization error
- Local force profiles are different, but overall forces are similar

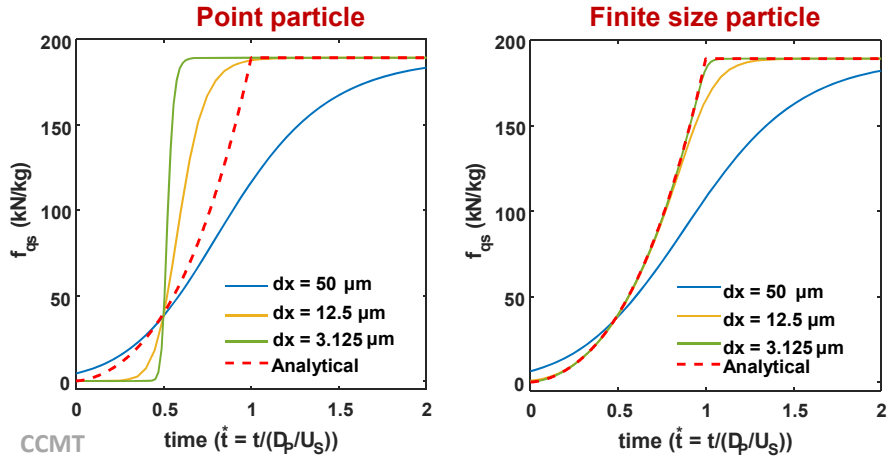

CCMT

21

Propagation of Error and Uncertainty

- Couplings of simulation and experiment for error estimation of the **key CMT physics** (**T2**,**T4**,**T5**,**T6**) of the demonstration simulation
- Applying **Forensic UQ** and **UR/ER** to the couplings for successful validation

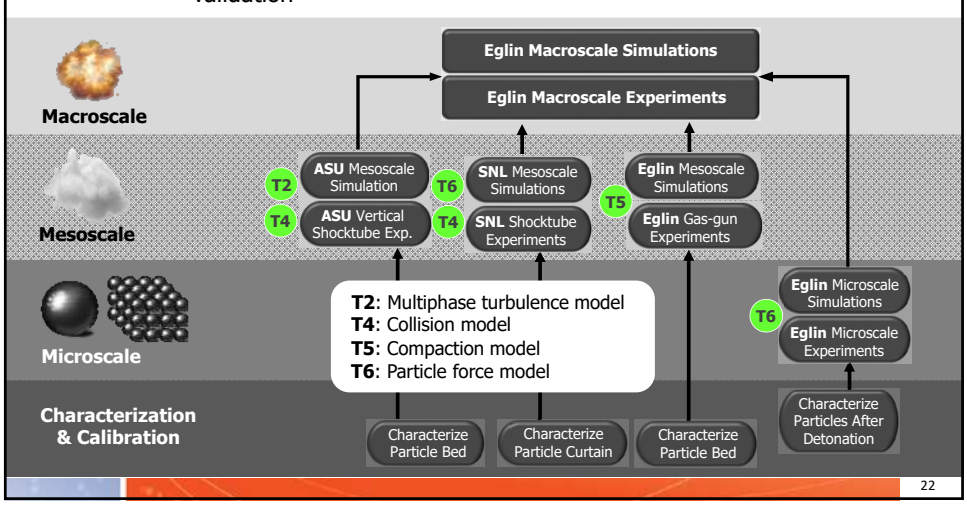

22

Propagation of Knowledge from Forensic UQ

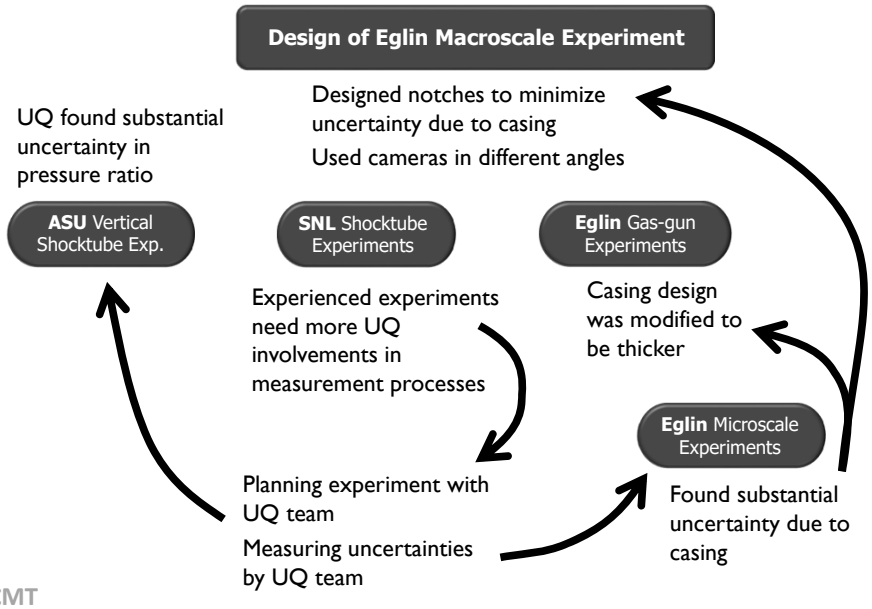

CCMT 23

Outline

- Unique aspects of UQ at CCMT
- UQ of demonstration problem
- VVUQ workflow of demonstration problem
- **Summary of Sandia Shocktube Campaign**
- Forensic UQ of micro- & meso-scale experiments

CCMT 24

Summary of Sandia Campaign

- Sandia campaign has reached useful failure after vigorous uncertainty reduction campaign
- Uncertainty in experimental data extraction process was significant
  - The uncertainty in data extraction process overwhelmed the experimental variability
  - Reduced uncertainty in data extraction because it was more critical than increasing repeatability

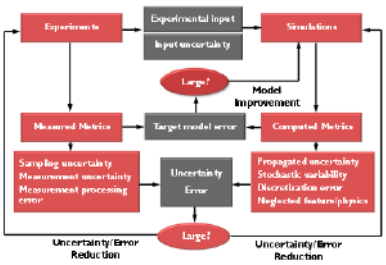

- Experiment configuration limited validation
  - Experiment plan need to consider simulation limitations
  - Actively involved in Eglin experiment plan
- **VVUQ framework** was developed and **is being used for the demonstration problem**

25

Shock Tube Experiment

- To observe shock-particle interactions in dense gas-solid flows
- Point particle force model is the key physics (Bertrand Rollin)
- Quantity of interest (QoI): Upstream and downstream particle front positions

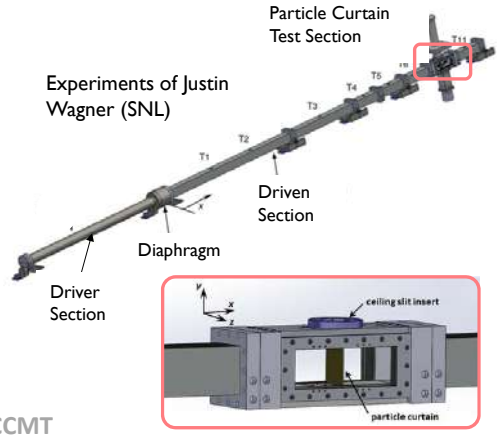

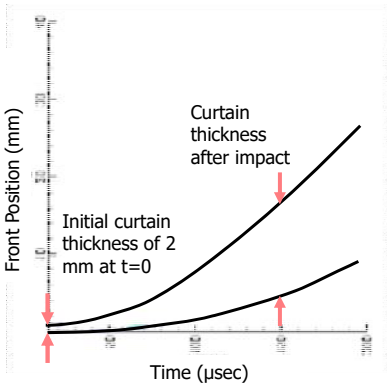

CCMT

26

Initial Uncertainties and QoI

- Key uncertainties (highest 6 out of 24)

Parameter	Method
Measurement bias in particle front positions	Uncertainty in systematic bias due to gap between particle curtain and wall
Particle volume fraction	Measurement uncertainty
Initial particle positions	Variability in initial particle positions
Particle diameters	Variability in particle diameters
Curtain thickness	Local curtain thickness variation
Pressure at driver section	Very small measurement noise
⋮	⋮

- Quantity of interest
  - Particle front positions **averaged** over initial particle positions
  - 4 repeated experiments
  - 10 repeated simulations with random initial particle positions

CCMT

27

Reducing Large Noisy Error due to AUSM+

Identified large error in AUSM+

Reduced the error by AUSM+up

ID

- Groups of runs identified large noise error in the behavior of average downstream front position for varying diameter **Forensic UQ**
- The physics team advanced the upgrade schedule of AUSM+ to AUSM+up **UR / ER**
- Model improvement substantially reduced the error

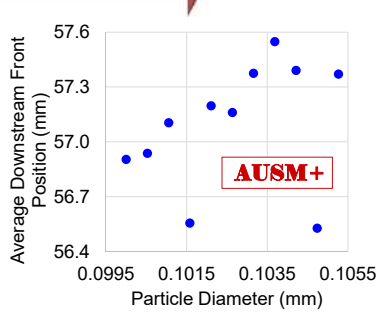

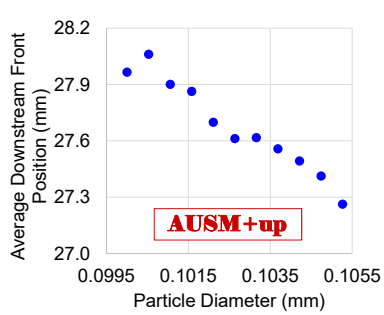

CCMT

28

AUSM+ vs. AUSM+up (1D)

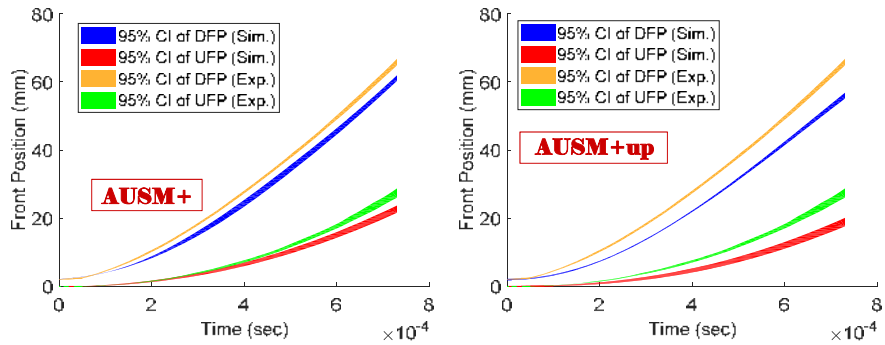

- Large noisy error due to AUSM+ was identified (Forensic UQ)
- Uncertainty reduction (UR) by implementing AUSM+up
- Discrepancy between simulation and experiment was significantly increased (discovered hidden error)

CCMT

29

Quantifying U. due to Gap Using 2D Sim.

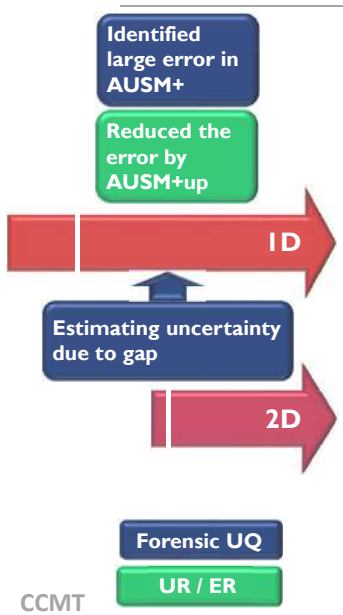

- 2D half particle curtain model was used to study the effect of gap
- Approximately 0-10% measurement bias in measured downstream front positions was expected

Biased estimation in 2D simulation

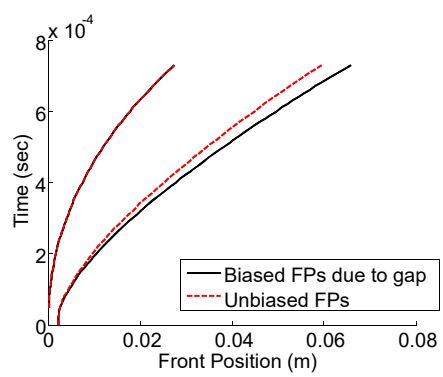

CCMT

30

Uncertainty Due to Gap Effect (1D)

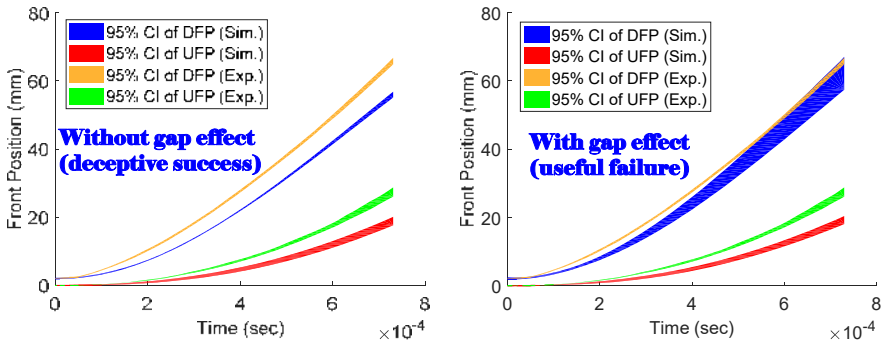

- The uncertainty due to gap should be considered in UQ (Forensic UQ)
- 2D simulation was used to estimate the uncertainty due to gap
- The second largest uncertainty was the uncertainty in the initial volume fraction (measurement uncertainty)

CCMT

31

Reducing U. due to Gap and U. in Initial VF

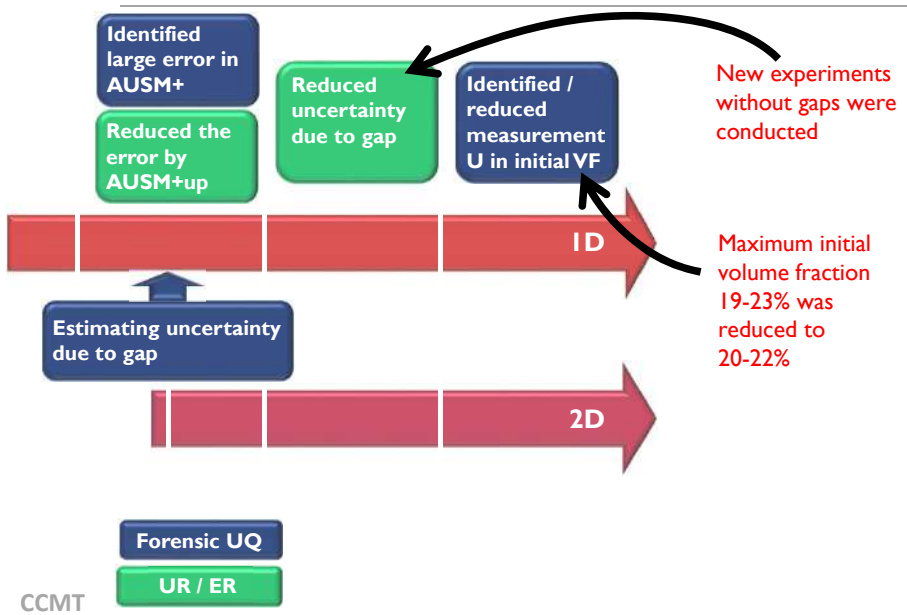

CCMT

32

1D and 2D After UR

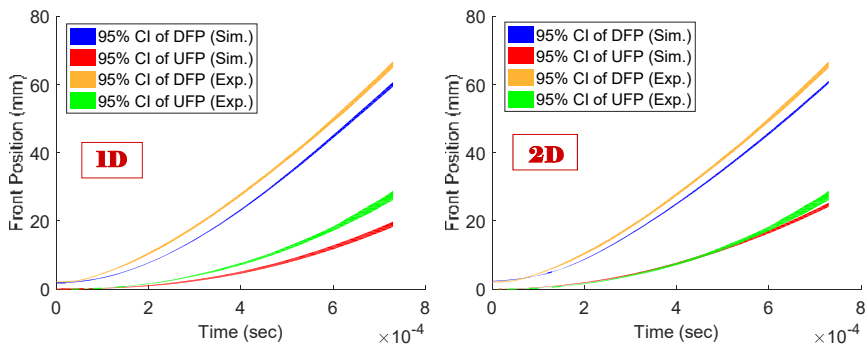

- The new experiment without gap removed the largest uncertainty due to gap was conducted
- The second largest measurement uncertainty in the initial volume fraction was reduced by rigorous UQ

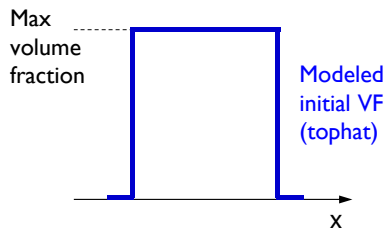

CCMT 33

Reducing Error in Initial VF Input

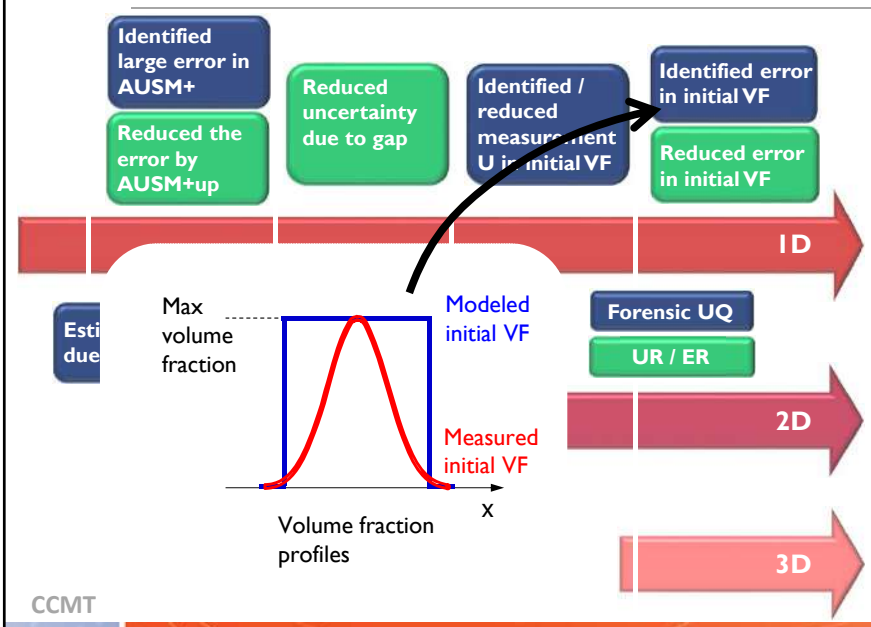

CCMT 34

1D/2D/3D After ER

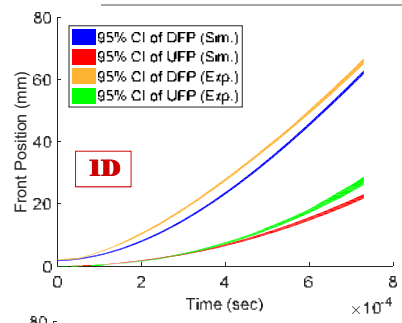

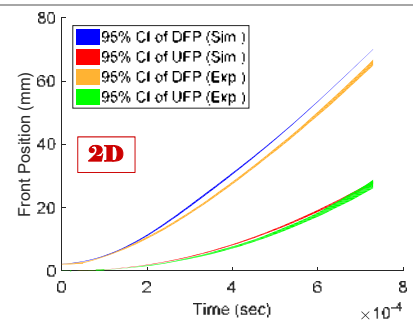

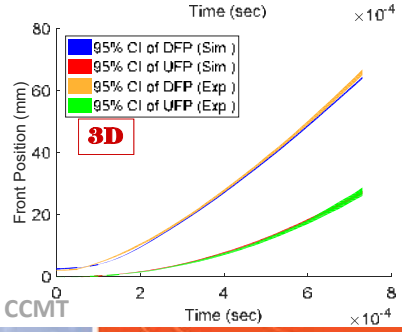

- Identified unrecognized error in the initial volume fraction input and was reduced (**Forensic UQ and ER**)
- The particle force model was applied in CMT-nek for ASU campaign
- Pseudo turbulence in 3D
- 81/625/6875 simulation runs for UQ of 3D/2D/1D

CCMT

35

Target Model Error Reduction Potentials

- Global sensitivity analysis (GSA) quantifies **target model error reduction potentials** of individual particle force models (Samaun Nili)
- GSA shows chances of large change of front position when one force model is improved
- Error in UFP is likely to be dominated by error in Quasi-Steady force model
- DFP was significantly influenced by errors in Quasi-Steady, Added Mass and Pressure Gradient force models

UFP

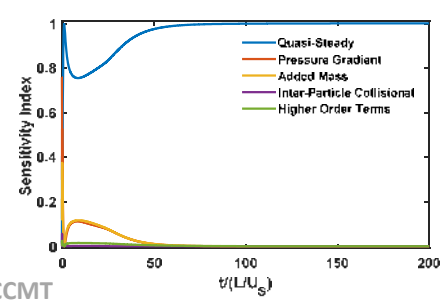

DFP

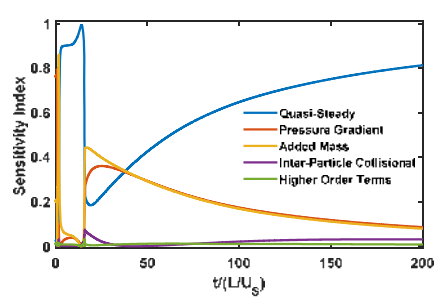

CCMT

36

Outline

- Unique aspects of UQ at CCMT
- UQ of demonstration problem
- VVUQ workflow of demonstration problem
- Summary of Sandia Shocktube Campaign
- **Forensic UQ of micro- & meso-scale experiments**

CCMT

37

Design of Macroscale Eglin Experiments

- 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

CCMT

38

Experimental Setup: February 2015

Side view photograph showing the concave pressure probe array.

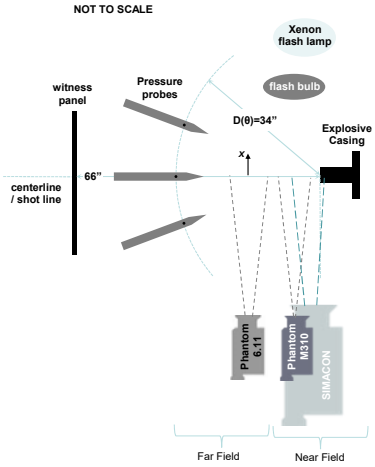

Overhead schematic of the test set-up.

Summary of test shots performed in October 2014 and February 2015. Particles were 2-mm tungsten spheres

Test #	Arrangement
Oct14-1	Single
Oct14-2	Ring of 7
Oct14-3	Grid of 16
Feb15-1	Single
Feb15-2	Single
Feb15-3	Diamond of 4

Sources: Myles Delcambre, Internal CCMT Fall 2014 presentation
Black, Littrell, and Delcambre, Eglin internal written report, 3/7/2015

CCMT

39

Major Uncertainty Sources: Casing Fracture

- Casing fragmentation is challenging to include in the simulation
- Designed new microscale explosives with a reduced amount of explosive and increased casing size (modeling after the mesoscale/p-rad)
  - 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)

Pictures of the microscale casing pre- and post- shot

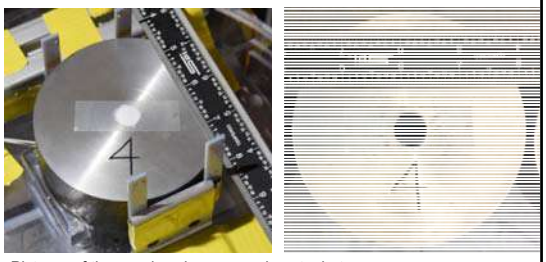

Pictures of the p-rad casing pre- and post- shot

CCMT

40

Microscale Gas Only Simulation

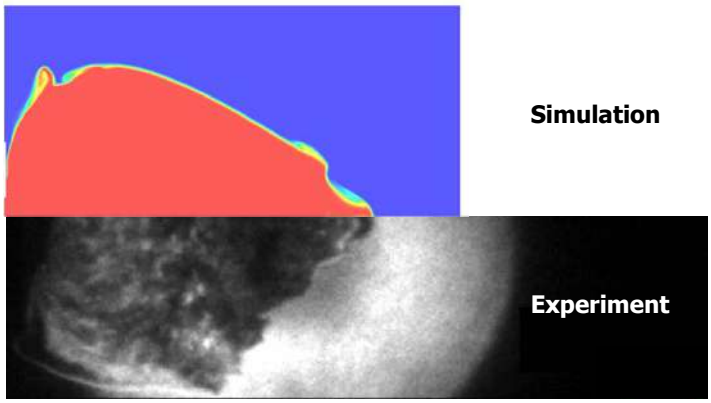

- Micro-scale gas-only simulation and experiment shows a good agreement
- This confirms that the error with a particle is solely due to the particle force model

CCMT

41

Error Estimation of Micro Sim. With One Particle

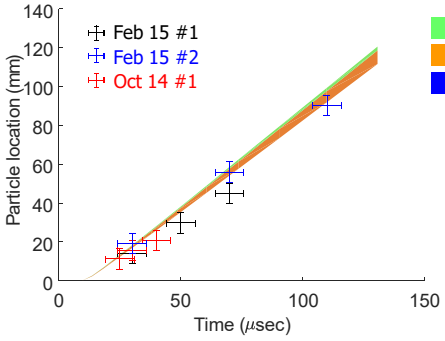

U in Density with CI 95%
U in Diameter with CI 95%
U in Explosive Density with CI 95%

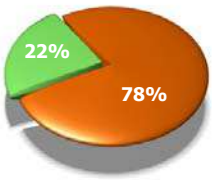

U in time (Exp)

U in location (Exp)

- The validated particle force model from the Sandia campaign is used
- The uncertainties in the inputs (particle diameter, density and explosive density) were propagated
- The uncertainty in the time of the experimental data will to be converted in terms of particle location for error estimation

CCMT

42

UQ in ASU Experiment

- **Goals:**
  - Plan new validation experiments that consider simulation limitations and uncertainty quantification goals
  - **Identify** and **quantify** sources of experimental uncertainty to have a greatest influence on simulation inputs
  - UQ is involved in the experiment (**Justin Mathew**, **Kyle Hughes** , **Tanner Jones**)
- Shock tube overview:
  - Four 32cm length segments
  - Particle diameter: 45-90 microns
  - Pressure ratio: 10~20
  - Initial bed height: 32 cm

43

Uncertain Experimental Parameters

Parameter	Quantity	Method
Driven section pressure	? kPa	Calculation based on shock speed
Driver section pressure	100.68±0 kPa	Static pressure sensor
Volume of particle bed	383.5-385.6 cm^3	
Initial volume fraction	59-65%	CT scanner
Initial mean volume fraction	60-61%	$\varphi_{mean} = \frac{m}{hA} * \frac{1}{\rho_{glass}}$ (Justin)
Particle diameter (before)	67±8 μm	SEM
Particle diameter (after)	68±22 μm	SEM
Particle density	2.44±0.006 g/cm^3	Pycnometer
Tube geometry	Negligible	

- **UB team** actively involved in measurements and uncertainty estimation (forensic UQ)
- The measured and calculated driven section pressure were different (forensic UQ)

CCMT

44

Particle Characterization – Shape/Size

- Glass particles: 45-90 microns in diameter, 2.44±0.006 g/cm^3 (**Kyle Hughes** and **Justin Mathew** at **UF**)
- Scanning electron microscope (SEM) Images show largely homogenous groups of particles with no apparent source of significant variation

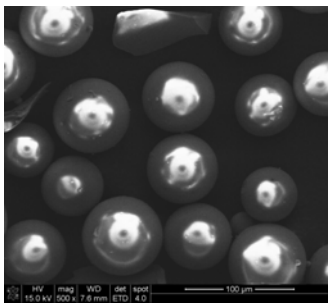

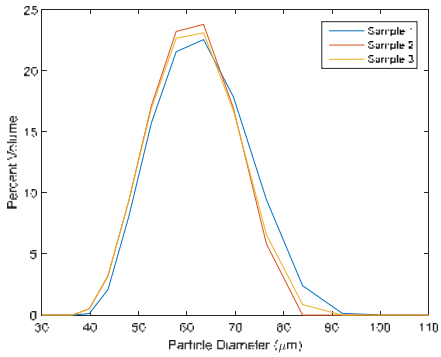

CCMT

45

Particle Bed Characterization – Initial VF

- To verify estimates of ~61% average volume fraction measured during ASU experiments, a mock particle bed was created for a **CT scan**
  - The mock tube was only 7 mm in diameter (versus ~40 mm shock tube diameter)
  - Observed a mysterious volume fraction increase in the radial direction

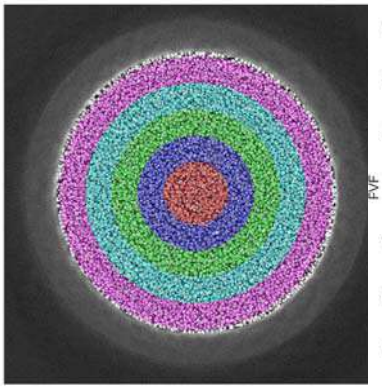

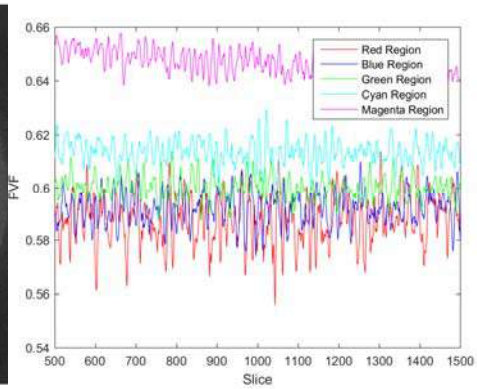

46

Comparison between Sim. and Exp. Results

- Mach number 1.494 (Oct/12 Exp #4)
- The height data from the experiments is based on high speed camera image that yields uncertainty in the measured height

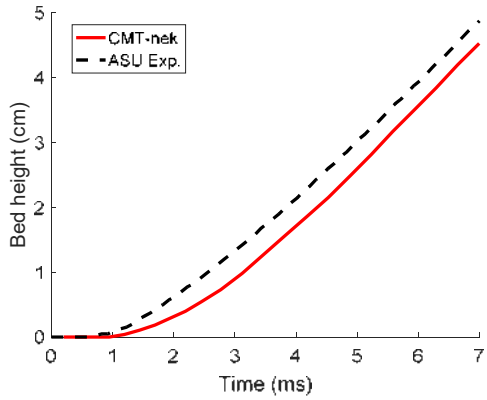

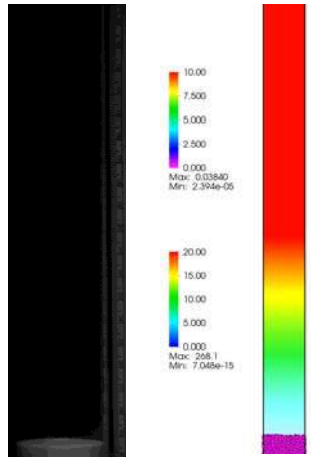

CCMT

47

Timeline

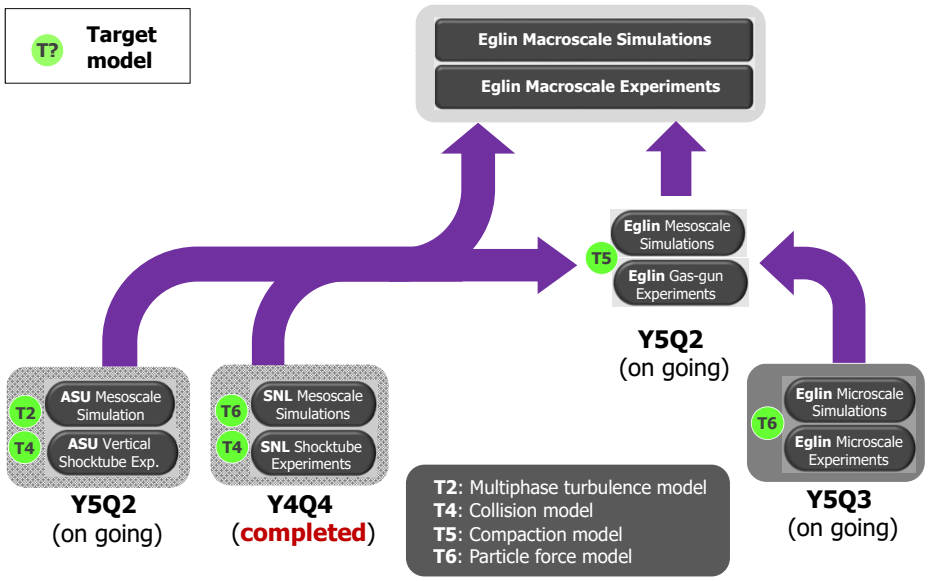

CCMT

48

Lessons Learned in Year 4

- Important to separate error from uncertainty for simulation development stage
  - Uncertainty reduction reveals simulation errors
- Forensic UQ has been fully exercised for different scales
  - Experience from micro- and meso-scale experiments were applied to macro-scale demonstration experiment
  - It is important that UQ team needs to be involved in experiment planning
- SNL shock tube simulation/experiment demonstrated that meaningful validation can be possible after reducing uncertainty
  - Based on 3D simulation, an important source of error is from 1D approximation

CCMT

49

Summary of Simulation Run Plan

- Groups of simulation runs (Mesoscale)
  - 1D/2D/3D simulations with a tophat initial VF profile for the new SNL experiments
  - 1D/2D/3D simulations with a bell-shaped initial VF profile for the new SNL experiments
  - Estimating discretization error of 1D/2D/3D simulations with finite volume particles
  - 1D/2D/3D simulations with finite volume particles and compare it to the new SNL experiment
- Groups of simulation runs (Mesoscale)
  - 2D/3D simulation (CMT-nek) and the ASU experiment (Justin and Chanyoung)
  - Axisymmetric mesoscale with a single particles, JWL and reactive burn
  - Axisymmetric gas gun with particles

CCMT

50

Summary of Simulation Run Plan

- Groups of simulation runs (Macroscale)
  - 1D macroscale with particles, JWL and reactive burn
  - 3D (or 2D) macroscale with particles, JWL and reactive burn with a case model

CCMT

51

CCMT

Do you have any questions?

Data Extraction from Shock Tube Experiment

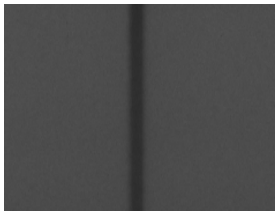

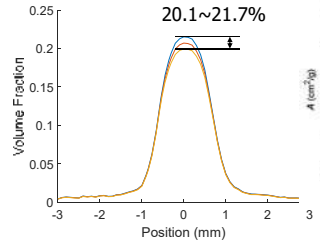

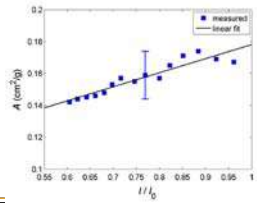

Calibration exp.

- X-ray to measure particle volume fraction profile
  - A processing model from X-ray image intensity to volume fraction
  - Need calibration for attenuation coefficient
  - Uncertainty in the volume fraction is from
    1) uncertainty in the experiment
    2) uncertainty in the calibration process (model form uncertainty)

Particle Curtain Edge Uncertainty

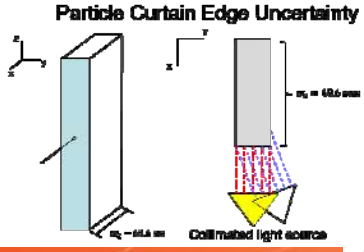

CCMT

53

ASU Experiments Forensic UQ

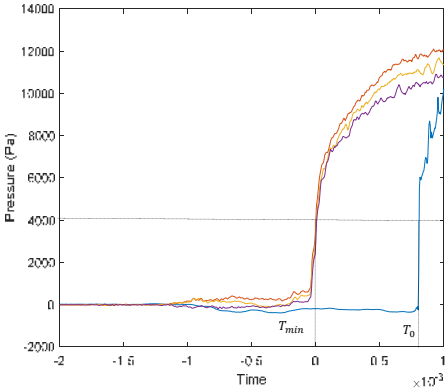

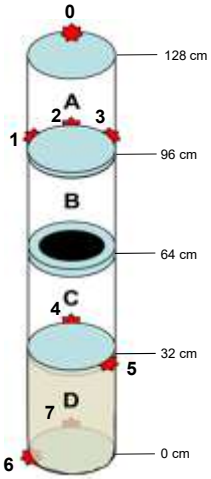

Experiment	$\frac{P_4}{P_1}$	Mach	Calc. $\frac{P_4}{P_1}$	% Error
0920_4	9.5	1.152	1.947	79.50
1012_5	13.3	1.393	4.851	63.52
1012_6	9.2	1.212	2.475	73.09

54

Data Extraction from Eglin Experiments

- Eglin microscale experiment (Forensic UQ)
  - Input UQ: particle diameter and density
  - Output UQ: particle positions and contact line

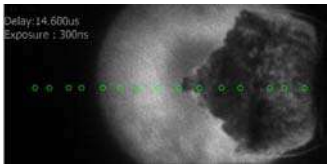

Simacon image of the Eglin microscale experiment

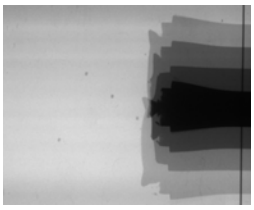

Multiple exposure x-ray of single particle

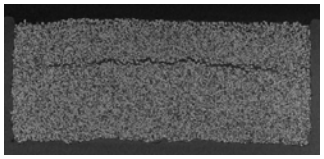

CT scans of particle packet for the Eglin mesoscale gas gun experiment

- Eglin mesoscale gas gun experiment
  - Computerized tomography (CT) scans of particle packet, SEM of particles, particle density via pycnometer
  - Simacon and X-ray images to track particle cloud movements

CCMT

55

CCMT

56

CCMT

CCMT
Student Lightning Round

UF UNIVERSITY of FLORIDA

DSE* of CMT-nek using BE**

Sai Prabhakar Rao Chenna
Advisor: Dr. Herman Lam, Dr. Greg Stitt
Department: ECE, UF

- Goals
  - Use BE methods & tools to perform extensive DSE on CMT-nek design space
  - Perform trace-driven simulations to optimize CMT-nek performance on HPC systems

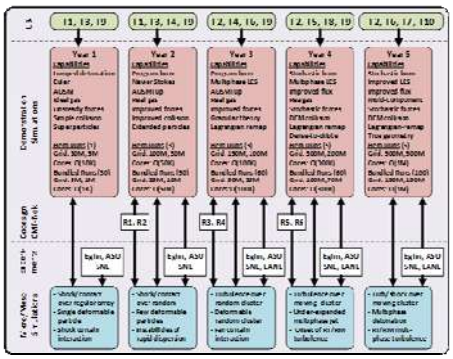

Simulation Roadmap

- Simulation roadmap
  - Early DSE to enhance CMT-nek performance on future exascale systems

* DSE - Design Space Exploration
**BE – Behavioral Emulation

CCMT

2

Uncertainty Quantification of Shock Tubes

Brad Durant

Advisor: S. Balachandar, T.L. Jackson

Department: MAE, UF

- Goals
  - Reproduce measurement processing from Sandia shock tube to identify sources of uncertainty
  - Propagate identified sources of uncertainty from experiments into simulations
- Simulation roadmap
  - T6: Force coupling
  - T4: Particle-particle interaction

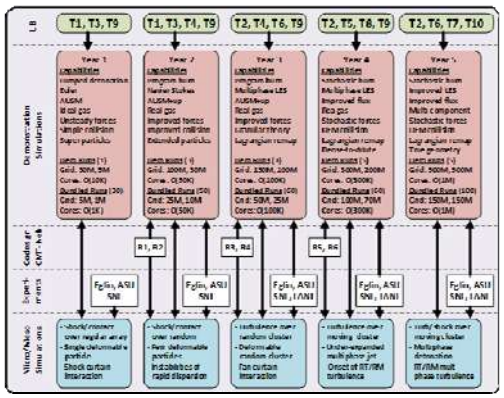

Simulation Roadmap

CCMT

3

Particles Departure from Axisymmetry

Maria Giselle Fernandez

Advisor: Dr. Raphael T. Haftka and Dr. S. Balachandar

Department: MAE, UF

- Goals
  - Determine what initial perturbation in particles leads to the maximum departure from axisymmetry
- Simulation roadmap
  - Collaboration with UB and Macroscale teams studying particles departure from axisymmetry

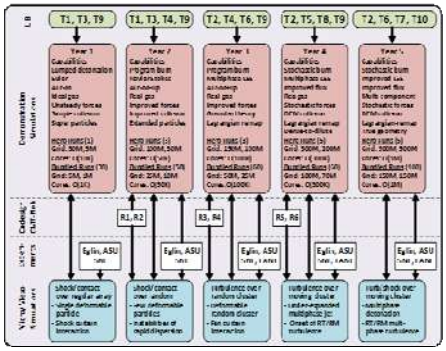

Simulation Roadmap

CCMT

4

CMT-bone on KNL and DVFS based Systems

Mohamed Gadou
Advisor: Prof. Sanjay Ranka
Department: CISE, UF

- Goals
  - Multiobjective Optimization of CMT-bone on DVFS based hybrid Systems
  - Performance optimization of CMT-bone on Intel KNL
- Simulation roadmap
  - Our research roadmap is closely tied with CMT-nek development.

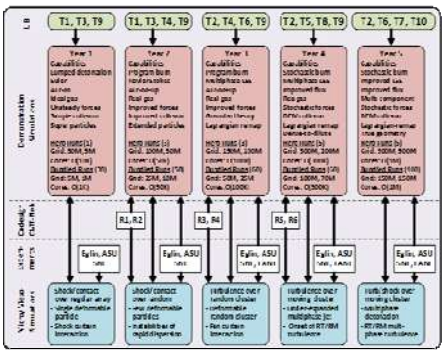

Simulation Roadmap

CCMT

5

Eglin Explosive Barrel Ejection Simulations

Joshua Garno
Advisor: Dr. S. Balachandar
Department: MAE, UF

- Goals
  - Simulate trajectory of 1, 9 and many particle configurations
  - Simulate dispersal of 0.5g packet of particles
  - Validate Rocflu models governing particle motion
- Simulation roadmap
  - Real Gas
  - EL-AUSM+UP
  - Improved Forces
  - Reactive Burn

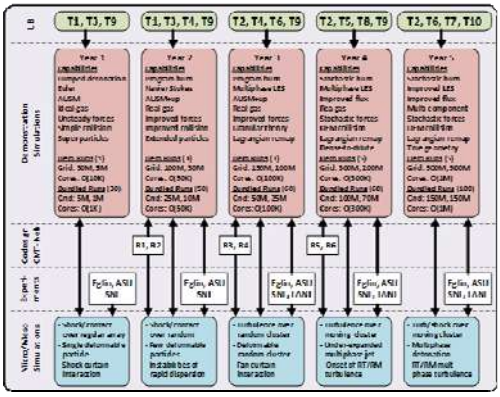

Simulation Roadmap

CCMT

6

Macroscale Explosive Dispersal of Particles

Kyle Hughes

Advisor: Prof. Kim, Balachandar, and Haftka

Department: MAE, UF

- Goals
  - Perform forensic uncertainty quantification of past experiments to inform future experiments
  - Design validation experiments to meet simulation assumptions and capabilities
  - Perform uncertainty quantification of validation experiments
- Simulation roadmap
  - Tie experiments performed at LANL and Eglin AFB to simulations using uncertainty quantification

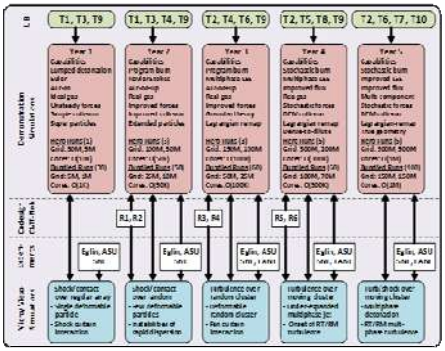

Simulation Roadmap

CCMT

7

Uncertainty Quantification of ASU Shock Tube

Tanner Jones

Advisor: Dr. Kim, Dr. Haftka, Dr. Park

Department: MAE, UF

- Goals
  - Identify sources of uncertainty in experimental procedures
  - Propagate identified sources of uncertainty from experiments into simulations
  - Validate CMT-nek simulation capabilities
- Simulation roadmap
  - Multiphase turbulence modeling and uncertainty
  - Conduct forensic UQ

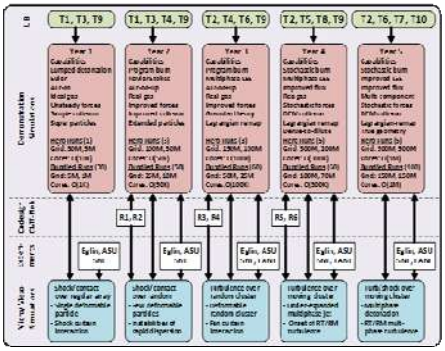

Simulation Roadmap

CCMT

8

Compressible Multiphase Turbulence Modeling

Rahul Babu Koneru
Advisor: Dr. S. Balachandar
Department: MAE, UF

- Goals
  - Improve drag models at high Mach numbers and dense regimes
  - Extend LES models to compressible multiphase flows
  - Validation and integration into CMT-nek
- Simulation roadmap
  - Validate existing force models and develop new models
  - Implement single phase LES and extend it to multiphase flows

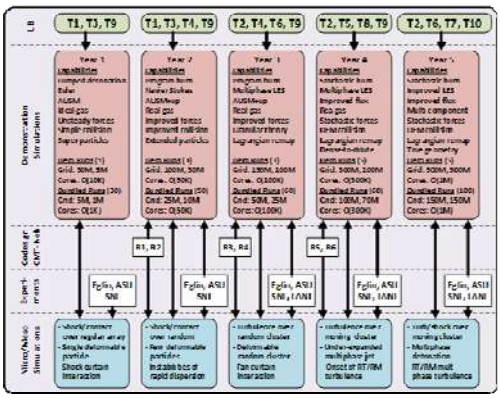

Simulation Roadmap

CCMT

9

CMT-nek Microscale Simulations

Goran Marjanovic
Advisor: Prof. Balachandar
Department: MAE, UF

- Goals
  - Demonstrate CMT-nek capabilities
  - Validation
  - Verification
- Simulation roadmap
  - Microscale simulations
    - Expansion wave propagating over particle beds

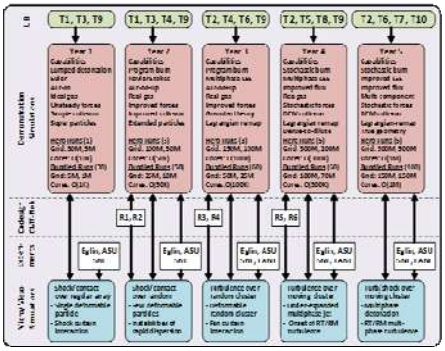

Simulation Roadmap

CCMT

10

Microscale – Shock Particle Interaction

Yash Mehta

Advisor: Prof. S. Balachandar

Department: MAE, UF

- Goals
  - Fully resolved DNS of shock interaction with particles
  - Developing models for predicting particle motion and force history
- Simulation roadmap
  - Simulating shock interaction with random bed of particles
  - Integration of force models in Meso-Macro scale simulations

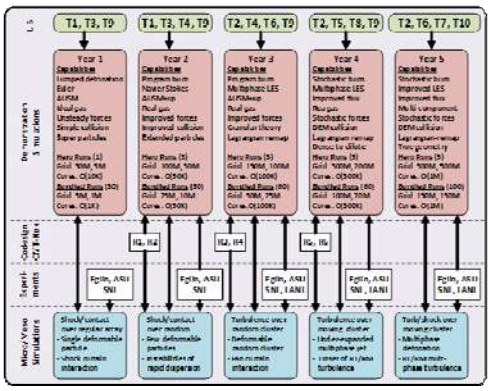

Simulation Roadmap

CCMT

11

Data-driven Point-Particle Force Modeling

W. Chandler Moore

Advisor: Prof. Balachandar

Department: MAE, UF

- Goals
  - Test the PIEP model at high volume fractions
  - Improve the PIEP model using DNS data
- Simulation roadmap
  - This work uses data from microscale simulations to improve force models which bridge the gap between microscale simulations and the demonstration simulations

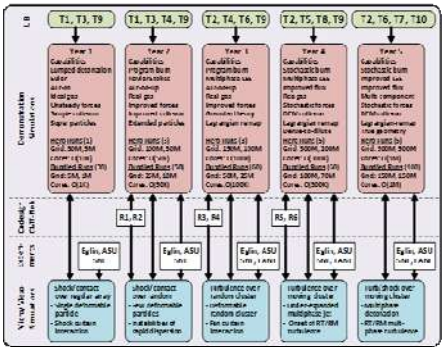

Simulation Roadmap

CCMT

12

Multi-fidelity Surrogate Modeling for Co-design

Aravind Neelakantan
Advisor: Dr. Lam, Dr. Stitt
Department: ECE, UF

- Goals
  - Reduce computational budget of HPC codes (parent app)
  - Low-cost model validation over larger design space
  - Extrapolation of parent app towards large scale run
- Simulation roadmap
  - MFS and BE helps in simulating CMT-nek on exascale systems (notional architectures)

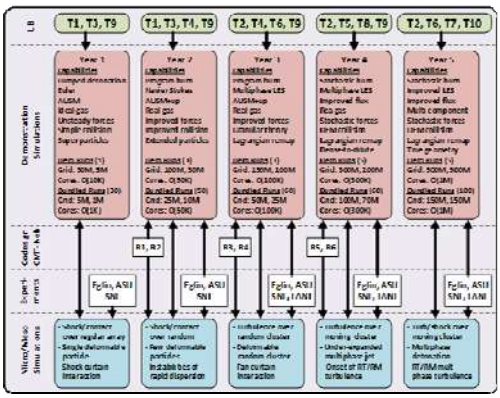

Simulation Roadmap

CCMT

13

Discretization Error in Euler-Lagrange Simulation

Samaun Nili
Advisor: Prof. Kim, Haftka & Balachandar
Department: MAE, UF

- Goals
  - To estimate the discretization error in a 1D Euler-Lagrange simulation of an interaction between the shock wave and particle-laden
  - To achieve mesh convergence for 1D shock particle-laden interaction problem via model implementation improvement
- Simulation roadmap
  - T4: Verification and validation of the shock tube simulation
  - Uncertainty reduction via model implementation improvement

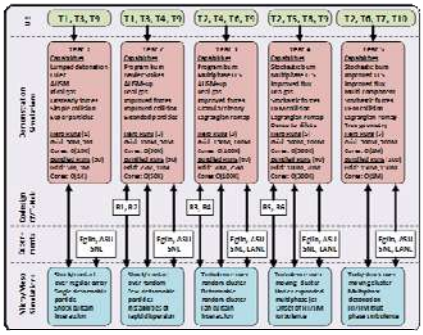

Simulation Roadmap

CCMT

14

Surrogate Modeling of the Equation of State

Frederick Ouellet
Advisor: Dr. S. Balachandar
Department: MAE, UF

- Goals
  - Develop surrogate models for use in evaluating equations of state (EoS) in mixed air/product cells
  - Perform and analyze macroscale simulations of the Frost and Eglin blastpad experiments
- Simulation roadmap
  - Real gas equation of state capabilities in code
  - Analysis of instabilities of rapid dispersion

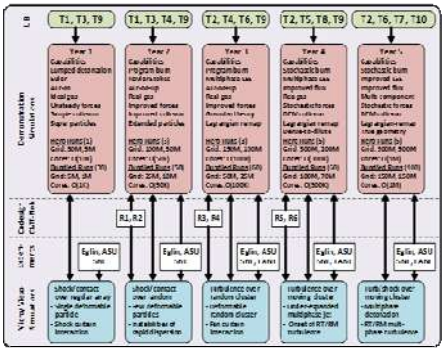

Simulation Roadmap

CCMT

15

FPGA Pipelined Simulations for CMT-nek

Carlo Pascoe, Ryan Blanchard
Advisors: Dr. Herman Lam, Dr. Greg Stitt

- Design-space exploration (DSE) critical in optimization of CMT-nek for potential Exascale architectures
  - CMT-nek has huge design space (DS)
  - BE is a potential solution, but is it enough?
  - Complement and accelerate BE approach via FPGA acceleration
  - Propose pipelined data flow simulations as useful technique for rapid exploration under certain circumstances

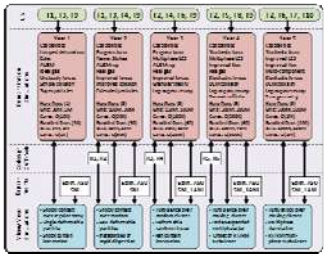

Simulation Roadmap

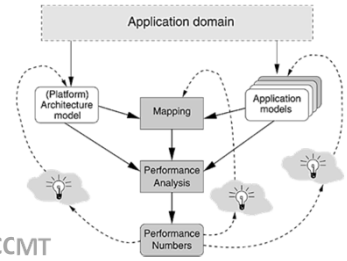

- Goal: Pipeline simulations to explore a different design option every cycle
  - After some initial latency, quickly explore large number of different configurations – fast DSE
- Ideal for an FPGA
  - Pipeline is bounded by resources rather than I/O
  - Multiple FPGAs can be used to explore a vast input space – almost embarrassingly parallel
  - **Up to eight orders-of-magnitude faster DSE**

CCMT

16

BE-SST Simulator

Rajashekar Rajagopalan

Trokon Johnson

Advisor: Dr Herman Lam,

Dr Greg Stitt

Department: ECE, UF

- Goals
  - Use SST framework to develop BE methods and run parallel, scalable simulations
  - Simulations and validations of large HPC systems
  - Improve communication models and Interpolation API
- Simulation roadmap
  - BE-SST simulator enables scalable Design Space Exploration on exascale systems

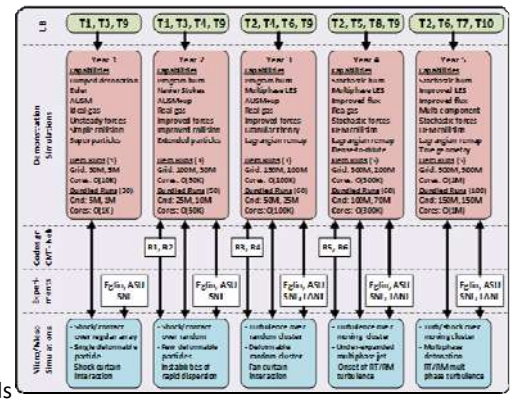

Simulation Roadmap

CCMT

17

Dynamic Load Balancing for CMT-nek

Keke Zhai

Advisor: Prof. Sanjay Ranka

Department: CISE, UF

- Goals
  - To reduce simulation time and decrease power consumption on CMT-nek by utilizing dynamic load balancing algorithms
- Simulation roadmap
  - This simulation includes moving particles and gas to simulate the actual particle movement within a fixed box

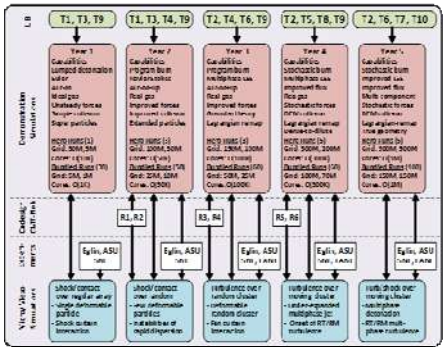

Simulation Roadmap

CCMT

18

UF UNIVERSITY of FLORIDA

Adaptive Sampling For Design Space Exploration Using Behavioral Emulation

Yiming Zhang
Advisor: Prof. Raphael T. Haftka
& Prof. Nam H. Kim
Department: MAE, UF

- Goals
  - Produce sampling plans and predictions using probabilistic models for the validation and usage of Behavioral Emulation of HPC architectures
- Simulation roadmap
  - Assist validation and usage of Behavioral Emulation for algorithmic & architectural design-space exploration at different stages

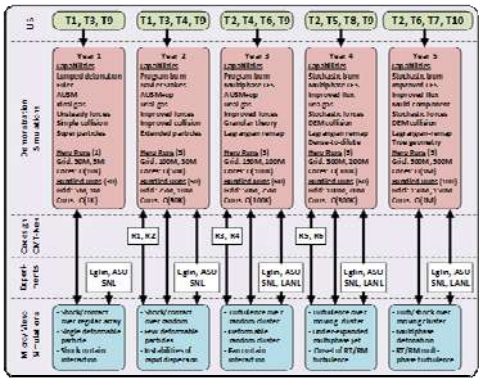

Simulation Roadmap

CCMT

19

UF UNIVERSITY of FLORIDA

Gas-Particle Mixtures Under Sudden Expansion

Heather Zunino
Advisor: Prof. Ronald Adrian
Department: SEMTE, ASU

- Goals
  - Perform repeatable experiments on a vertical shocktube at ASU
  - Examine expansion fan, flow structures, and instabilities
  - Provide data for validation of computational codes
- Simulation roadmap
  - The shocktube experiment at ASU will provide data for validating the codes being developed at UF

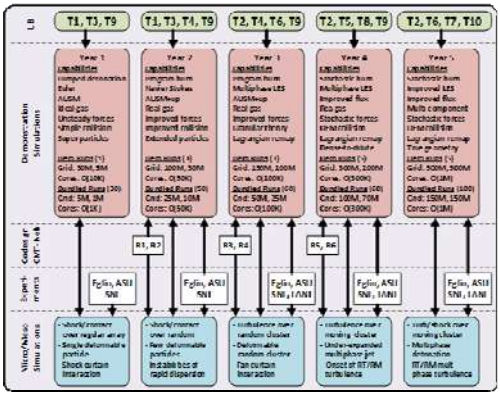

Simulation Roadmap

CCMT

20

Simulations of Fine Powder Decompression

David Zwick
Advisor: Prof. Balachandar
Department: MAE, UF

- Goals
  - Development of state-of-the-art Eulerian-Lagrangian capabilities in CMT-nek
  - Understanding the physics of particles interacting with expansion waves
- Simulation roadmap
  - Detailed modeling of particle simulations of ASU experiment in CMT-nek

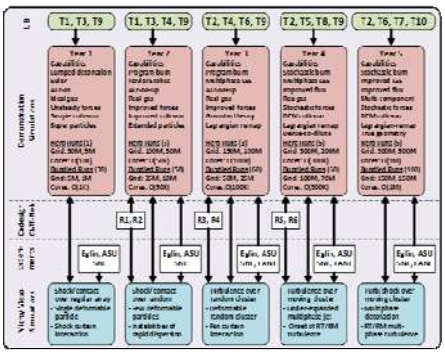

Simulation Roadmap

CCMT

21

CCMT

The End