Accelerating predictive simulation of IC engines with high performance computing (ACE017)

K. Dean Edwards (PI), Charles E. A. Finney, Wael R. Elwasif
National Transportation Research Center
Oak Ridge National Laboratory

DOE EERE Sponsors:
Gurpreet Singh, Michael Weismiller
Vehicle Technologies Office

2019 DOE Vehicle Technologies Office Annual Merit Review
10-13 June 2019; Arlington VA USA

Overview

Timeline

- Project start – FY2012
- Current efforts under FY2017-2019 Lab Call
- Proposed to continue through FY2023 under Combustion Consortium

Barriers

- Targets key barriers from USDRIVE roadmap for modeling and accelerating development of advanced engines with improved efficiency and emissions
  - "… develop more robust, computationally efficient models for combustion system design for improved efficiency and reduced CO2 emission. The knowledge base and modeling tools are required to design combustion systems for maximum fuel economy and minimum emissions." – 2018 ACEC Roadmap

Budget

- FY2017 – $340k
- FY2018 – $400k
- FY2019 – $400k

Partners

- Leveraging DOE Office of Science ASCR leadership computing resources
  - Multiple ALCC and DD allocation awards
  - Currently totaling **~20 Mhrs** on Titan & other OLCF resources
- Supports collaborative efforts with industry through OLCF ACCEL industrial partnership program
- Collaboration with LLNL for application of GPU-enabled Zero-RK chemistry solvers on Titan: **ACE012**
- Strong connectivity to broader simulation portfolios at ORNL and DOE
  - Co-Optima simulation team: **FT071**
  - Daimler SuperTruck II team: **ACE100**
  - Advanced propulsion materials: **MAT057**

Supporting industrial partnerships through the OLCF

Project supports collaborative efforts with industry, other NLs, and academia using HPC to improve simulation capabilities enabling accelerated development of advanced engine designs

Oak Ridge Leadership Computing Facility (OLCF)

Oak Ridge National Laboratory | Leadership Computing Facility

- DOE User Facility funded by the Office of Science through the Advanced Scientific Computing Research (ASCR) program
- Yearly competitive proposal process for resource allocations: DD, ALCC, INCITE*

Titan, 27 PF
18,688 nodes: 16 CPU, 1 GPU
32 GB/node

Summit, 200 PF
4608 nodes: 42 CPU, 6 GPU
600 GB/node

ACCEL
Accelerating Competitiveness through Computational ExceLence

ACCEL: OLCF Industrial Partnership Program

- Provides path for industry to collaborate with ORNL researchers on complex problems that require HPC resources
- Open, precompetitive (early TRL) projects (inputs may remain proprietary)
- Expectation that results will be published in open literature

GM, FCA, Ford, GE + Oak Ridge National Laboratory | National Transportation Research Center

- Supported tasks address technical barriers that are…
  - Of particular interest to the industry partner… and transportation industry as a whole
  - Within scope of DOE-VTO research focus areas
  - Well-suited for the use of DOE's leadership HPC resources
  - Precompetitive (early TRL) with results publishable in open literature
- Multi-year efforts must be coordinated through yearly OLCF project allocations

* DD = Director's Discretion (typically 1-5 Mhrs)
ALCC = ASCR Leadership Computing Challenge (5-20 Mhrs)
INCITE = Innovative and Novel Computational Impact on Theory and Experiment (20+ Mhrs)

Relevance and *Approach*

Using HPC and GPUs to "move the bar" on the speed-accuracy trade-off

- HPC and GPU-enabled solvers enable ***higher-fidelity*** CFD simulations and ***faster throughput***
  - *Accuracy:* More detailed physics- and chemistry-based submodels
  - *Speed:* Increased throughput with parallel simulation
  - *Speed:* Reduced computation time with GPU-based solvers
- Today's supercomputers are tomorrow's workstations
  - GPU-based systems increasingly available and affordable

Methods and techniques developed on HPC now will be ready for future commercial systems

GM, ORNL, LLNL, and CSI are partnering on multi-year effort to apply this approach to virtual engine calibration and design

Unique tools…

High-performance computing

OAK RIDGE National Laboratory

+

GPU-enable kinetics solver

Zero-RK | Lawrence Livermore National Laboratory

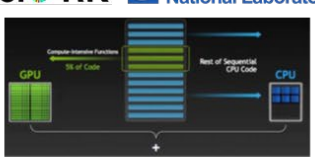

enabling high-fidelity simulation…

Detailed CFD engine model

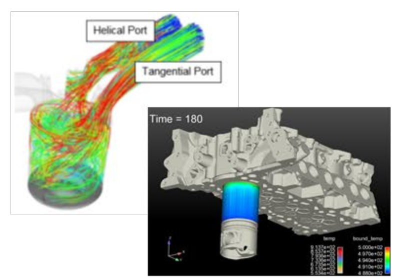

CONVERGE CFD SOFTWARE | CONVERGENT SCIENCE

for Industry applications

Virtual engine calibration & design

GM General Motors

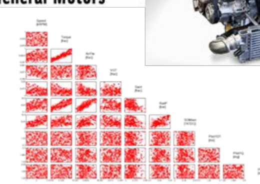

OAK RIDGE National Laboratory | NATIONAL TRANSPORTATION RESEARCH CENTER

Relevance and *Approach*

Virtual calibration has potential to significantly accelerate engine development

- Use high-fidelity, predictive simulations to generate data for the calibration process
  - Reduce need for hardware testing
- <u>Barrier</u>: *speed vs. accuracy trade-off*
  - Requires simulation of 1000s of calibration points
  - Includes extreme operating conditions that are challenging to simulate
    - Full-authority sweeps of ECU-level control parameters, rich operation, high EGR, etc.
- For conventional computing resources, need simple models for acceptable throughput
  - ***Can often match combustion performance with acceptable accuracy…***
  - ***but <u>emissions</u> predictions needed for calibration development are significant challenge***

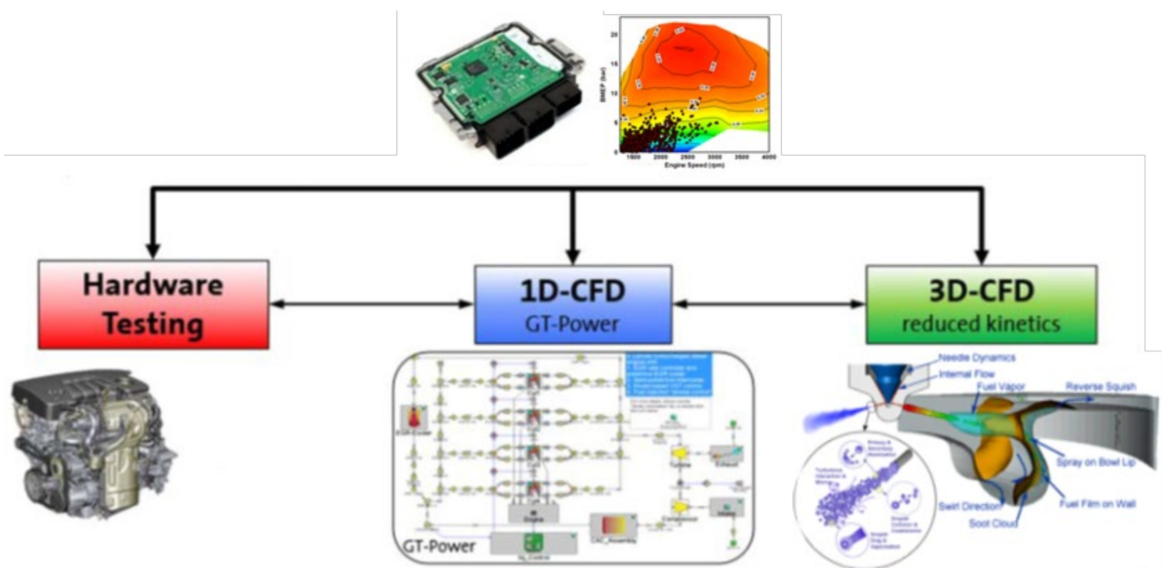

Relevance and Approach

Engine platform and experimental validation data

- GM 1.6-L mid-size Diesel
- Down-selected 602 experimental calibration points for simulation
  - Represents DoE of 10 ECU-level controls
  - Includes extreme cases difficult to capture with single model

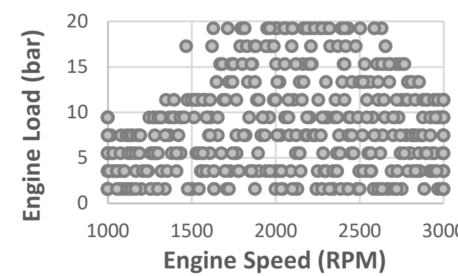

Goals:

- Improve accuracy
- Increase prediction (less tuning)
- Increase coverage of operating space
- Prioritize model refinements
- Approach applicable to both CI and SI

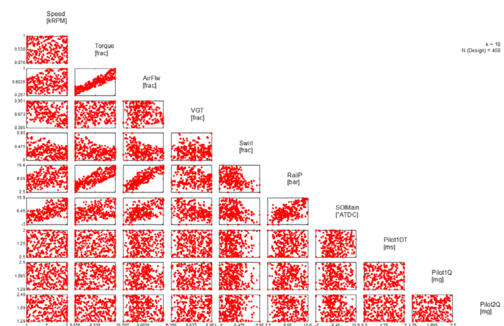

Systematic approach to add detail to CFD model

- Detailed chemical kinetics w/ GPUs
- Mesh refinement
- Spray model refinement

2015 (CI) & 2016 (SI) ALCC

- Sector vs. full-cylinder model
  - Closed-valve vs. open-cycle with gas exchange

2017 ALCC

- Conjugate heat transfer
  - Full multi-cylinder engine model
- LES turbulence model

2018 ALCC

Multi-year effort performed through yearly ASCR project allocations at OLCF

Overview of past results

Baseline model represents simplified approach typically required for adequate throughput on conventional computing resources

	Baseline
Geometry	Sector
Cycle	Closed
Max # cells	~170k
# species	47
# reactions	74
NOx	Zel'dovich
Soot	Detailed PSM
Turbulence	RANS
Other changes	
Wall time / cycle	~2 hr
Cases	500

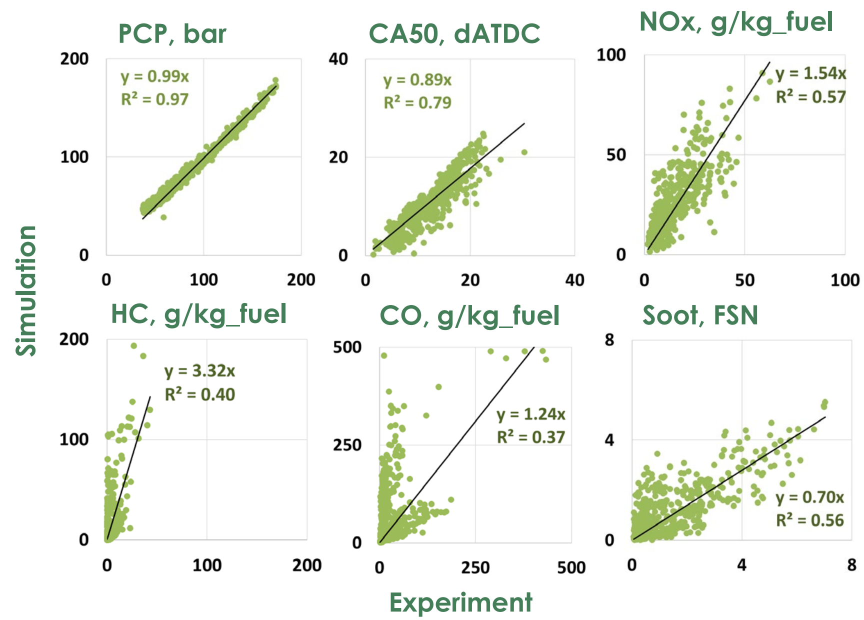

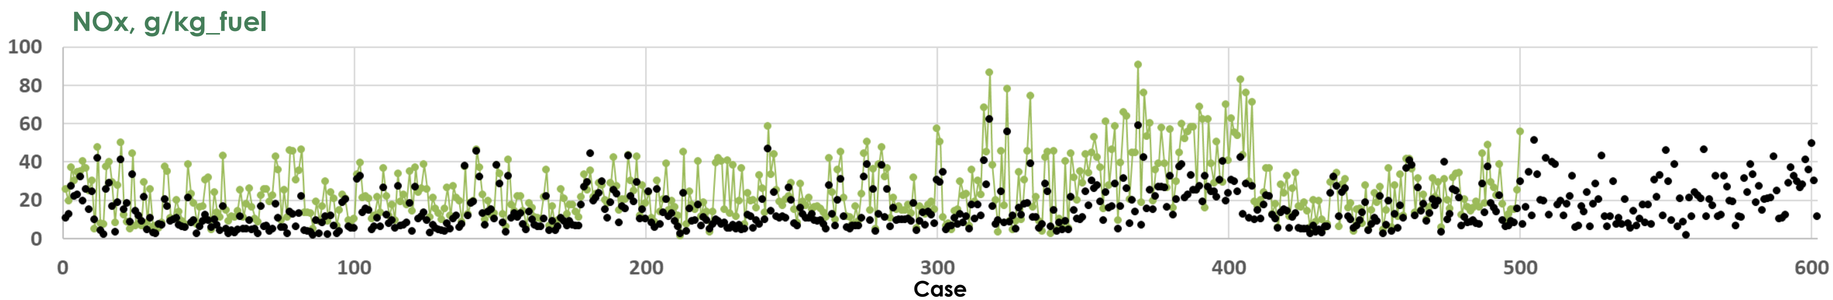

Adequately captures combustion metrics but performs poorly for most emissions

Overview of past results

Best sector-model results obtained with refinements to chemistry, mesh, spray, and wall-film

	Baseline	Best Sector
Geometry	Sector	Sector
Cycle	Closed	Closed
Max # cells	~170k	~1M
# species	47	144
# reactions	74	900
NOx	Zel'dovich	GRI 3.0
Soot	Detailed PSM	Hiroyasu
Turbulence	RANS	RANS
Other changes		Spray and wall-film
Wall time / cycle	~2 hr	~5 hr
Cases	500	602

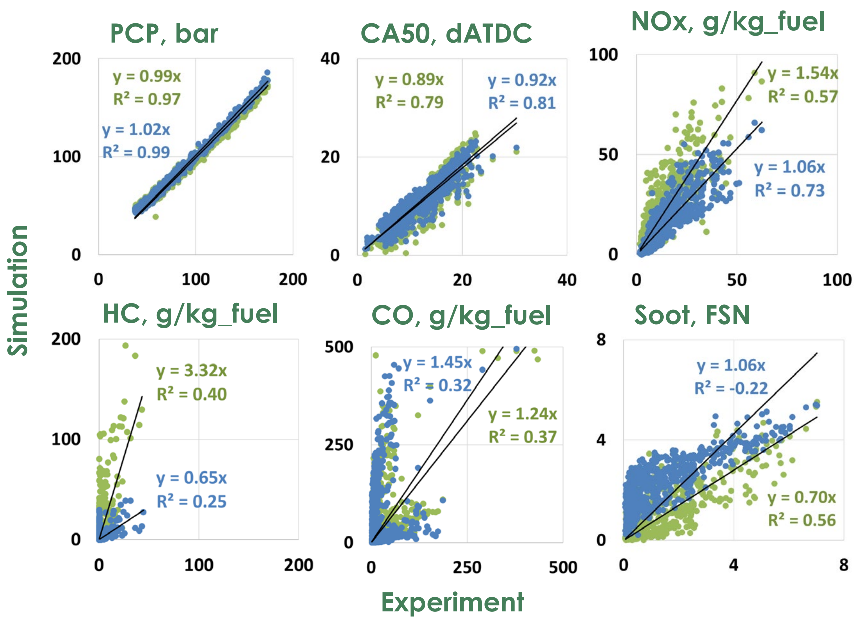

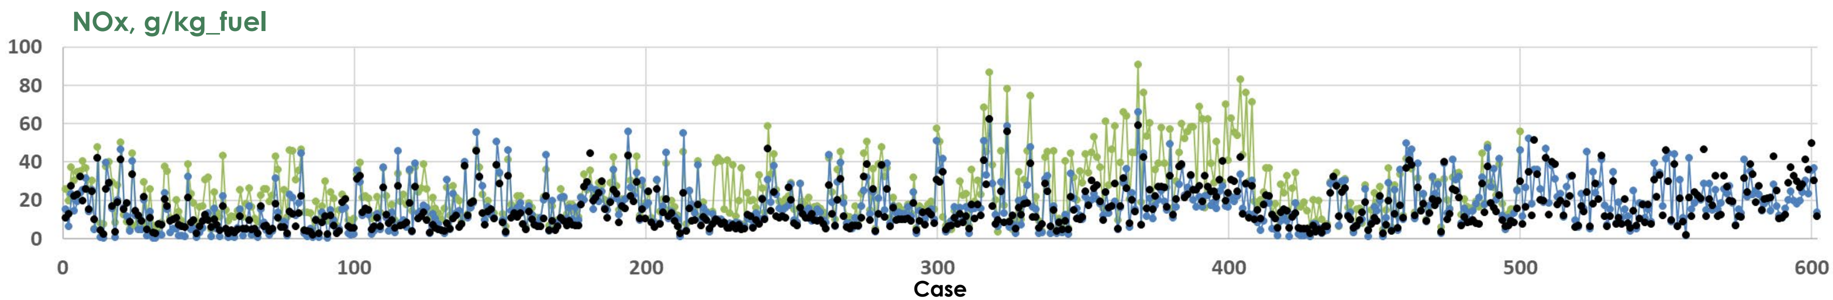

NOx: Significant improvement with refinement to chemistry and grid (and thus temperature)

Overview of past results

Best sector-model results obtained with refinements to chemistry, mesh, spray, and wall-film

	Baseline	Best Sector
Geometry	Sector	Sector
Cycle	Closed	Closed
Max # cells	~170k	~1M
# species	47	144
# reactions	74	900
NOx	Zel'dovich	GRI 3.0
Soot	Detailed PSM	Hiroyasu
Turbulence	RANS	RANS
Other changes		Spray and wall-film
Wall time / cycle	~2 hr	~5 hr
Cases	500	602

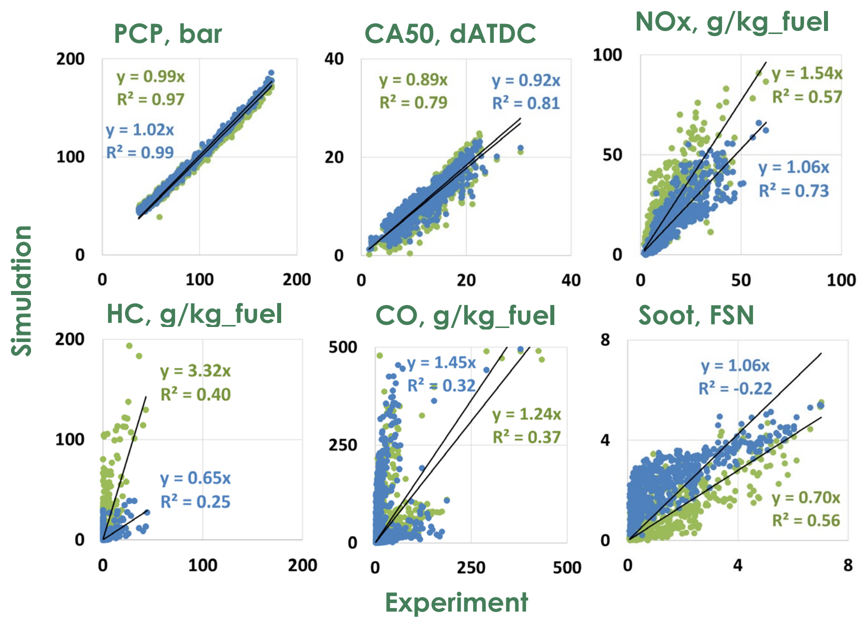

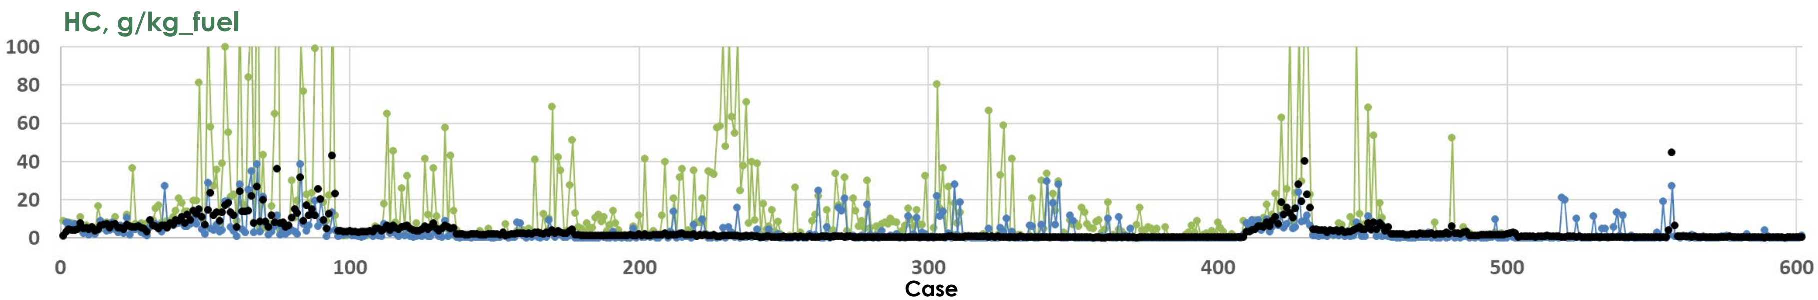

HC: Significant improvement with chemistry and wall-spray refinements

Overview of past results

Best sector-model results obtained with refinements to chemistry, mesh, spray, and wall-film

	Baseline	Best Sector
Geometry	Sector	Sector
Cycle	Closed	Closed
Max # cells	~170k	~1M
# species	47	144
# reactions	74	900
NOx	Zel'dovich	GRI 3.0
Soot	Detailed PSM	Hiroyasu
Turbulence	RANS	RANS
Other changes		Spray and wall-film
Wall time / cycle	~2 hr	~5 hr
Cases	500	602

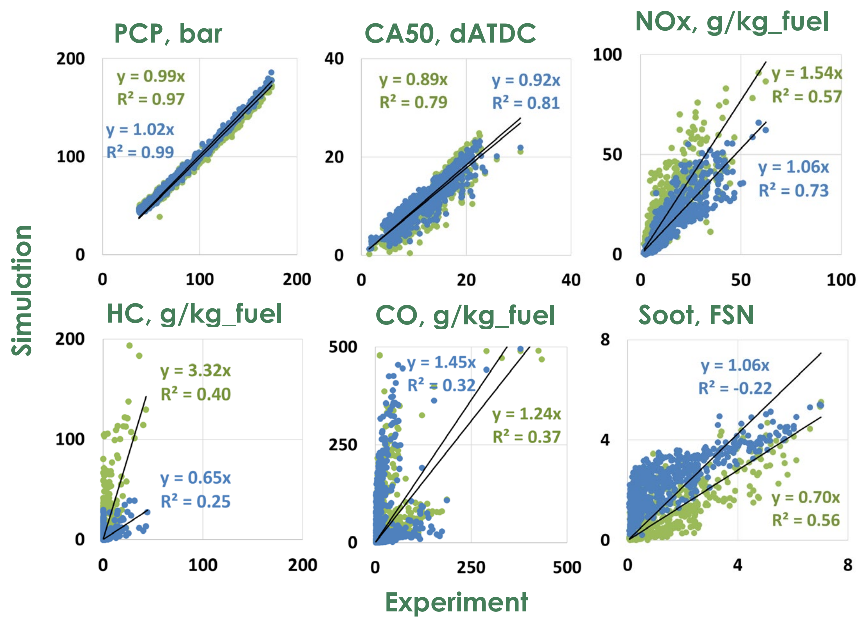

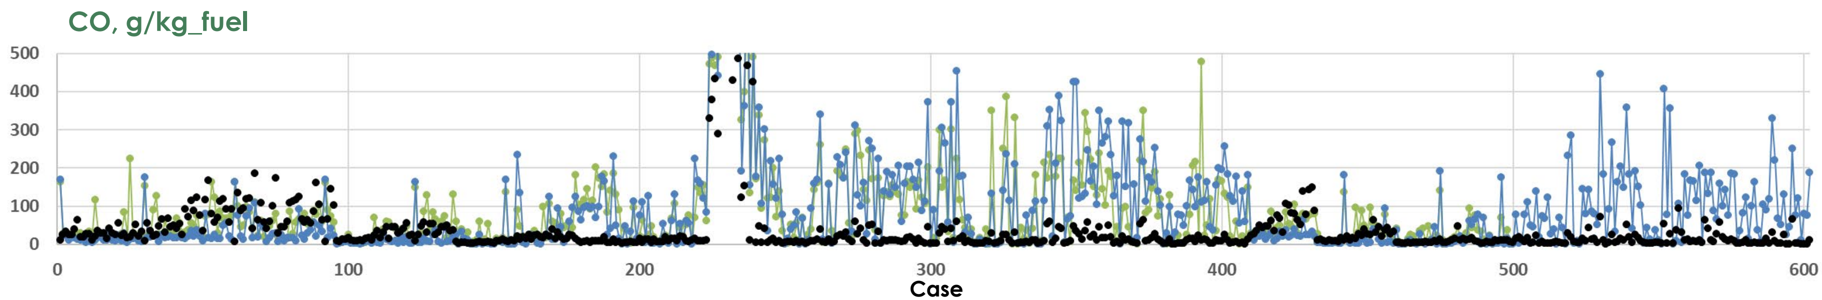

CO and soot: Need further work

Overview of past results

Open-cycle with full-cylinder geometry

	Baseline	Best Sector	Full-cyl
Geometry	Sector	Sector	Cylinder
Cycle	Closed	Closed	Open
Max # cells	~170k	~1M	~2.5M
# species	47	144	144
# reactions	74	900	900
NOx	Zel'dovich	GRI 3.0	GRI 3.0
Soot	Detailed PSM	Hiroyasu	Hiroyasu
Turbulence	RANS	RANS	RANS
Other changes		Spray and wall-film	Intake swirl vane
Wall time / cycle	~2 hr	~5 hr	~3.5 days
Cases	500	602	20

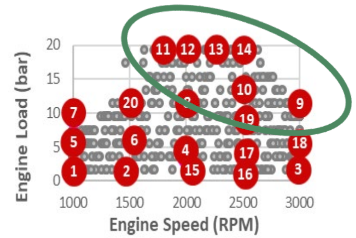

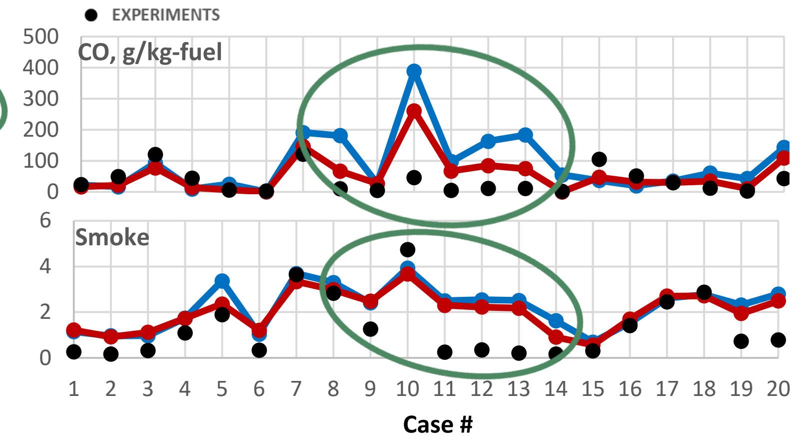

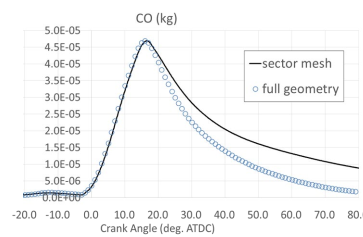

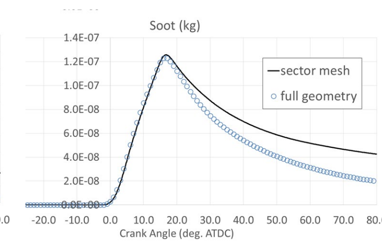

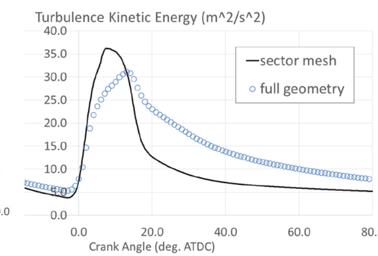

- 20 down-selected cases (due to allocation limits)
- Notable improvements in CO and soot
  - Further improvement needed at high speed and load
- Higher late-cycle oxidation rates
- Earlier, more complete CO oxidation in squish region
- Improved small-scale mixing

LES is expected to further resolve mixing effects

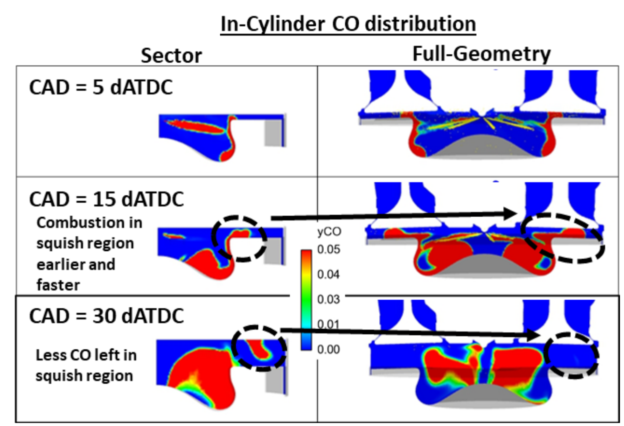

Adding CHT for spatio-temporal thermal BCs and LES

Thermal BCs are difficult to measure and poorly known

- Often assume constant, uniform wall temperatures from estimations or 1-D thermal solvers (*e.g.*, GT-Power)
- CHT moves BC from cylinder walls to locations where boundaries are better defined

Multi-cylinder engine model developed with conjugate heat transfer (CHT) to solve for thermal BCs

- CONVERGE v2.4.20 with 3-D CHT
- Compound model structure with iterative solution
- Combustion solved in cylinder 1 with thermal solution mapped to other cylinders

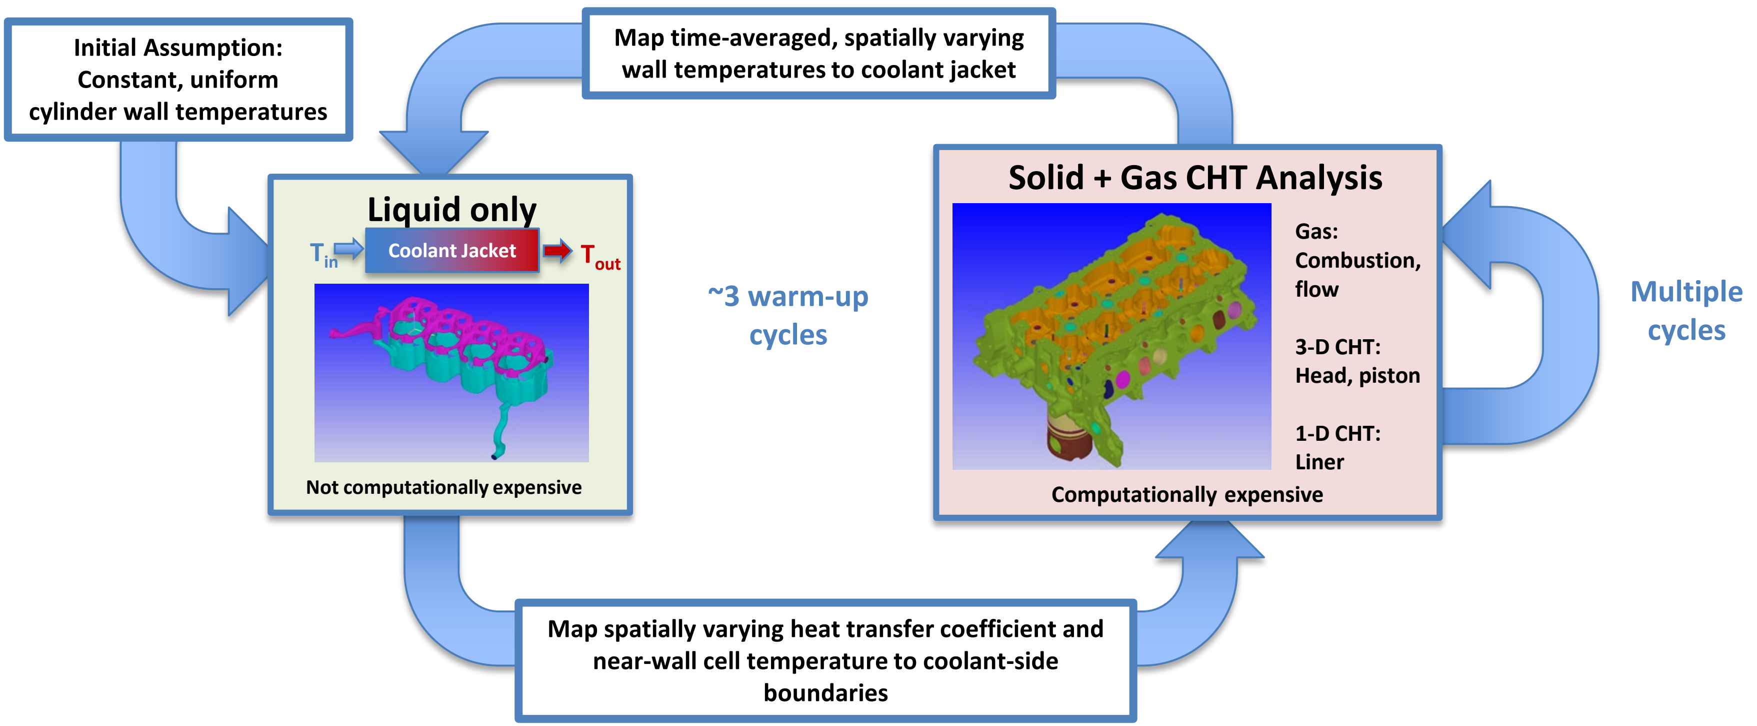

Technical Accomplishments

CHT runs with RANS and LES are underway on Titan

	Baseline	Best Sector	Full-cyl	CHT+RANS	CHT+LES
Geometry	Sector	Sector	Cylinder	Engine	Engine
Cycle	Closed	Closed	Open	Open	Open
Max # cells	~170k	~1M	~2.5M	~4.5M	~5M
# species	47	144	144	144	144
# reactions	74	900	900	900	900
NOx	Zel'dovich	GRI 3.0	GRI 3.0	GRI 3.0	GRI 3.0
Soot	Detailed PSM	Hiroyasu	Hiroyasu	Hiroyasu	Hiroyasu
Turbulence	RANS	RANS	RANS	RANS	LES
Other changes		Spray and wall-film	Intake swirl vane	CHT	CHT
Wall time / cycle	~2 hr	~5 hr	~3.5 days	~2 weeks	~3 weeks
Cases	500	602	20	8	8

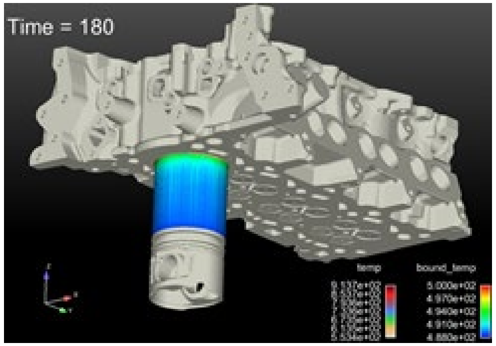

Full multi-cylinder engine geometry

- 3-D CHT on piston and head
- 1-D CHT on liner
- Combustion solved in one cylinder
  - Thermal conditions mapped to other cylinders

- Further down-selected to 8 cases (due to allocation limits)
- 3 warm-up cycles with imposed HR
- Multiple simulations branched from end of warm-up
  - 3-4 cycles with **RANS**, improved mechanism (Zero-RK)
  - "X" cycles with **LES**, improved mechanism (Zero-RK)
    - Where X = as many cycles as possible with remaining allocation
  - 1 case for 5 cycles with **RANS, skeletal** mechanism (SAGE)

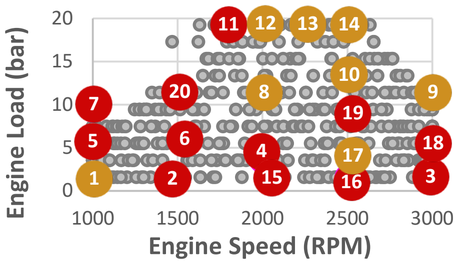

Technical Accomplishments

- Preliminary results from combustion cycles with RANS, skeletal mechanism
  - 2500 RPM, 19 bar

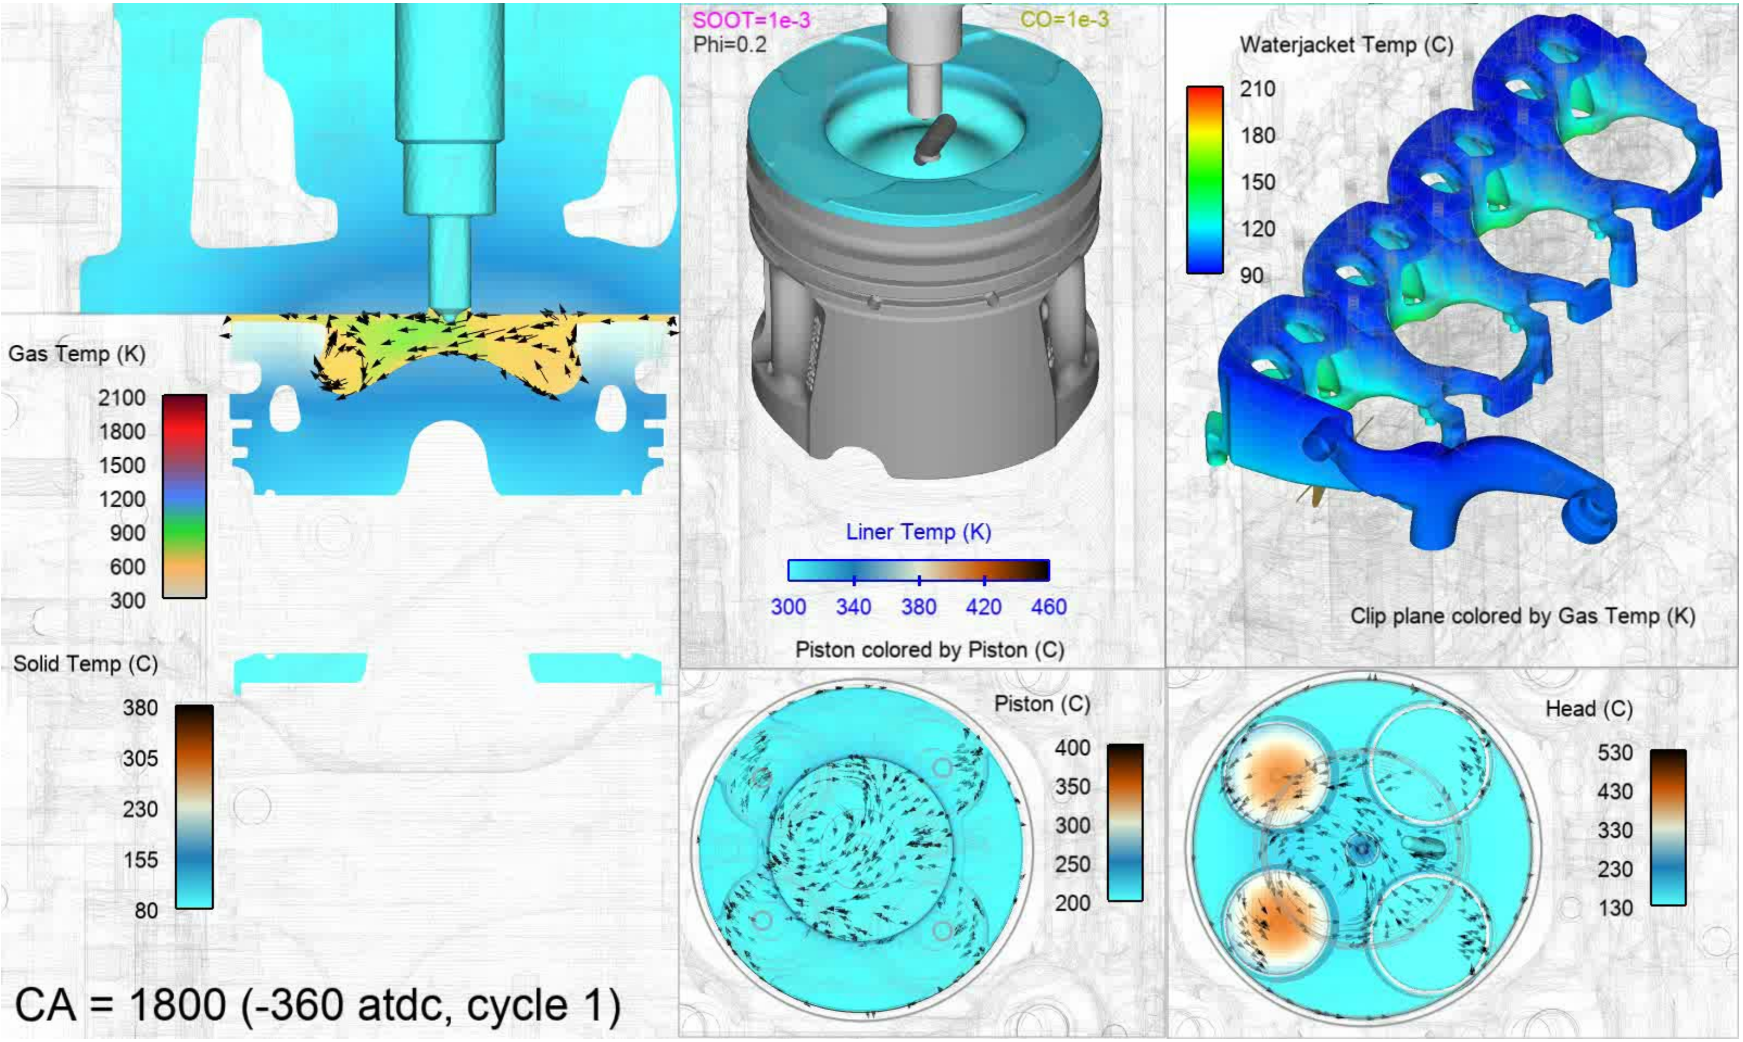

Oak Ridge National Laboratory | National Transportation Research Center

Technical Accomplishments

Initial results: Charge motion comparison for RANS

- Swirl ratio (bulk motion) similar for sector and full-geometry models
  - Noticeably higher for CHT cases (note scale)
  - Both RANS and LES show similar results
  - Possible thermal effect
- Small-scale turbulence predictions similar during injection for both geometries
- Turbulent KE higher and length scale larger during expansion for full-geometry model

Increased small-scale motion during expansion coincides with observed higher CO and soot oxidation rates

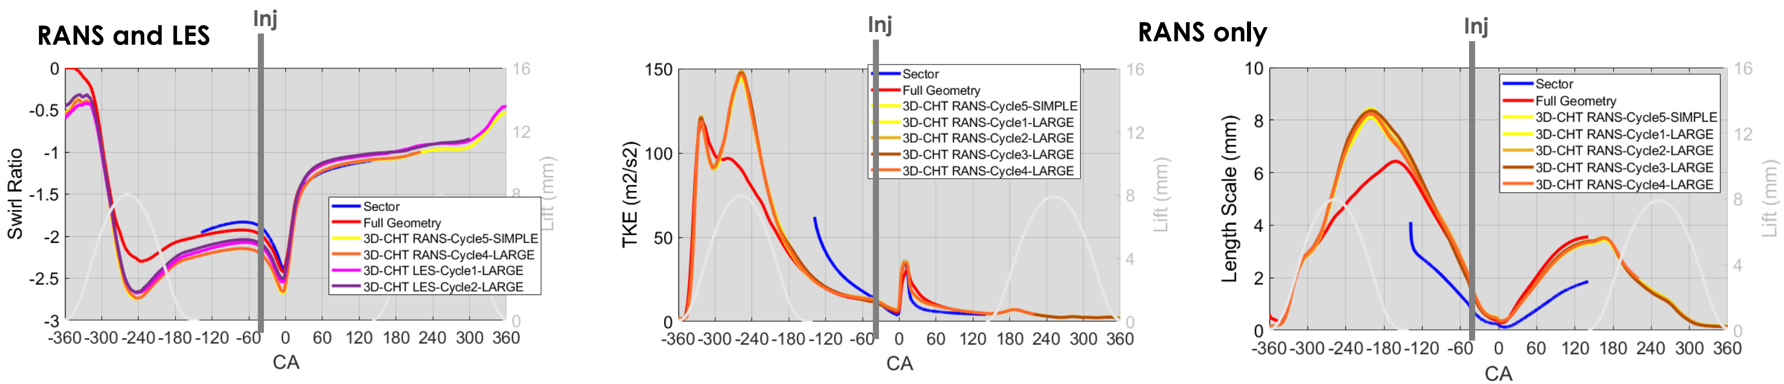

NOTE: Results shown are from a single case (2500 RPM, 19 bar)

Technical Accomplishments

Initial results: Thermal BC comparison

- CHT+RANS solution converges in 2-3 cycles (after 3-cycle warm-up)
  - Note: 1-D CHT on liner does not use supercycling so still converging
- CHT+RANS predicting significantly different spatially-averaged wall temperatures than 1-D GT-Power calibration used for sector and full-geometry models
  - Piston: ~60°C higher   Head: ~30°C lower   Liner: similar
- Moderate temperature swing (temporal variation) predicted with CHT+RANS
  - Piston: ~30-40°C   Head: <10°C   Liner: <10°C

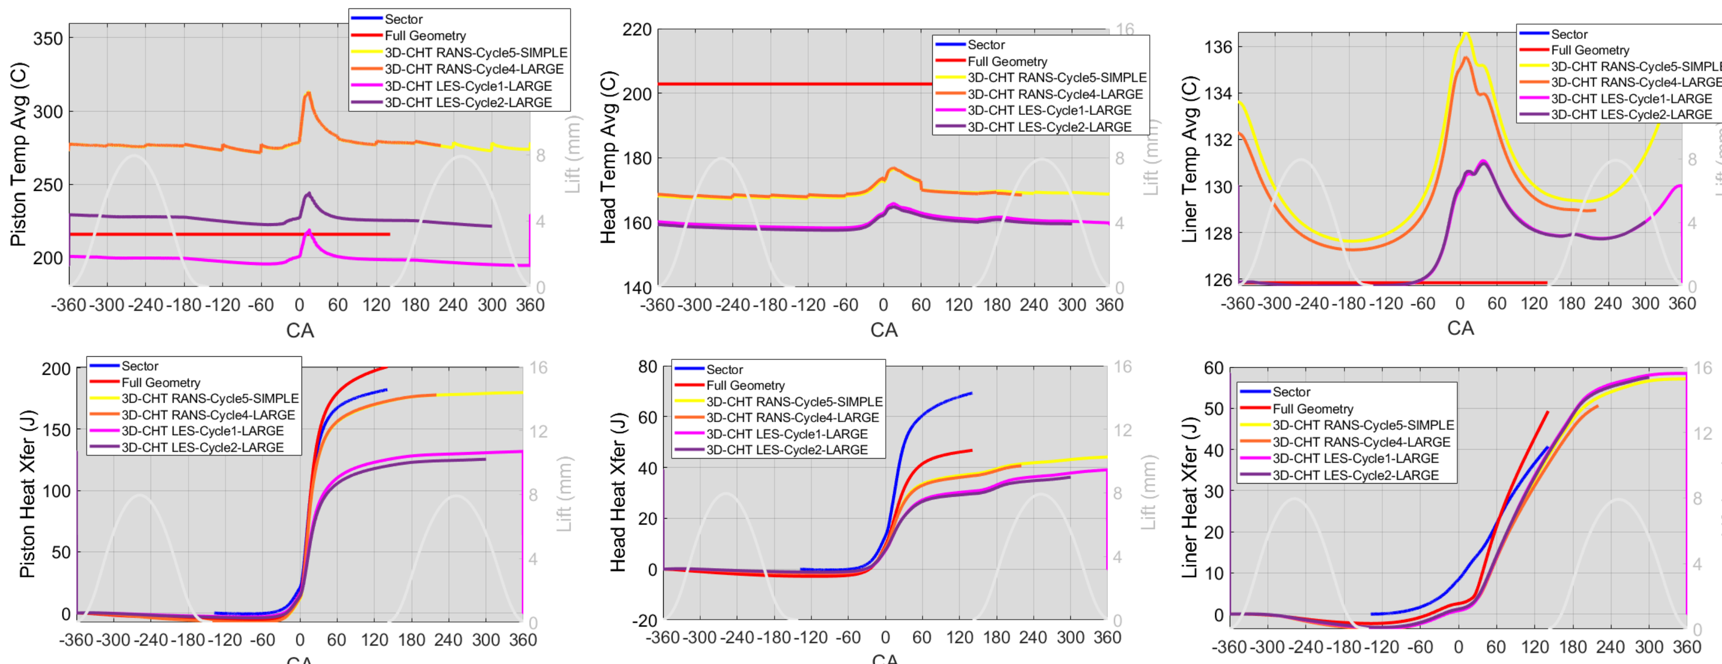

NOTE: Results shown are from a single case (2500 RPM, 19 bar)

Technical Accomplishments

Initial results: Thermal BC comparison

- CHT+LES solution appears to be converging more slowly
  - Piston temperature increase on 2nd cycle much less than for RANS
  - Head and liner temperatures show little difference between 1st and 2nd cycles
- CHT+LES also predicting lower temperature swings (*e.g.*, ~20°C for piston vs. 30-40°C for RANS)
- Lower heat transfer losses from piston and head with CHT+LES despite cooler wall temperatures

Initial results suggest that RANS is predicting higher thermal diffusion rates than LES

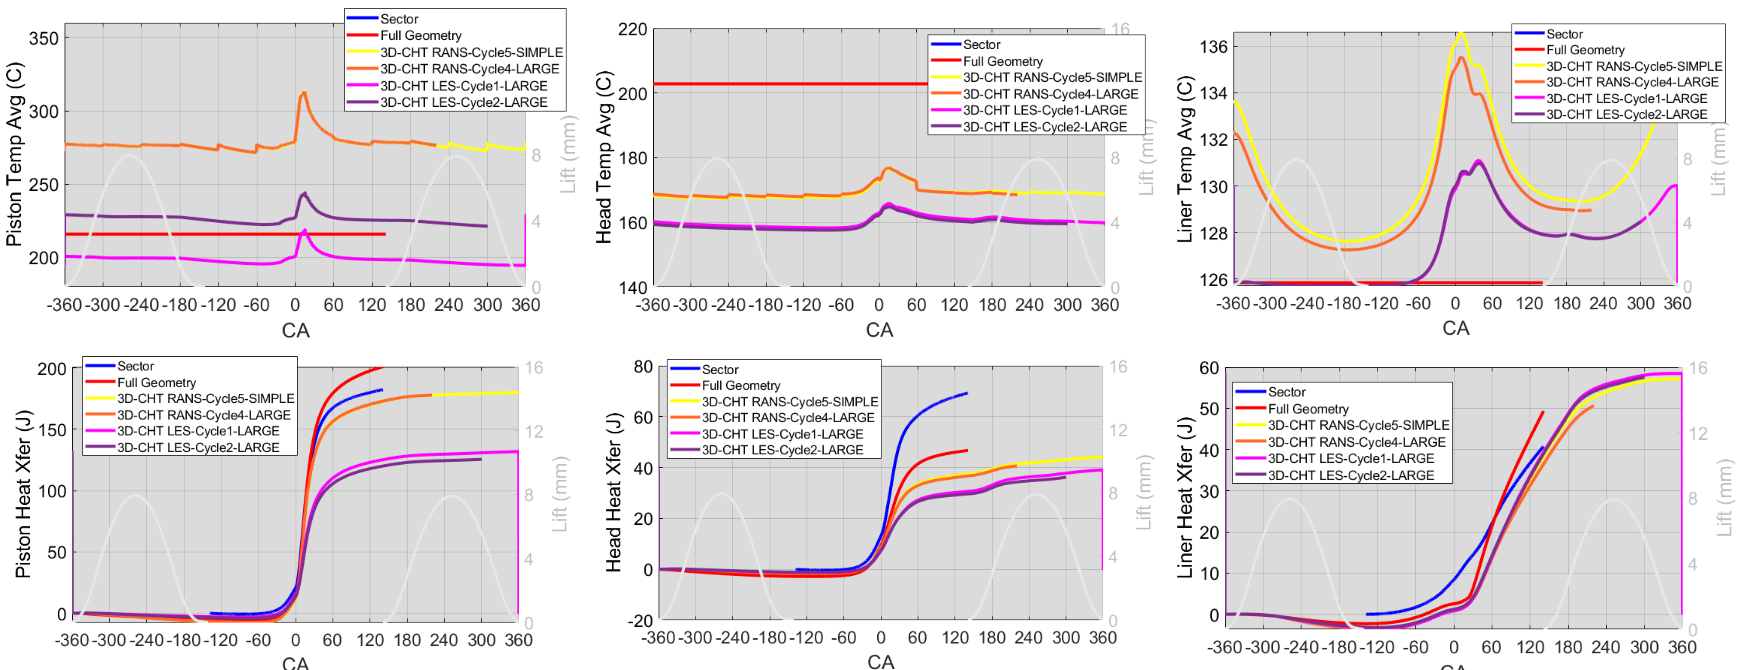

NOTE: Results shown are from a single case (2500 RPM, 19 bar)

Technical Accomplishments

Initial results: Thermal BC comparison

- Considerable spatial variation in piston (130°C) and head (80°C) temperatures
  - Piston temperature from 1-D GT-Power model matches minimum predicted with 3-D CHT+RANS
  - Head temperature from 1-D GT-Power model matches maximum predicted with 3-D CHT+RANS

Piston

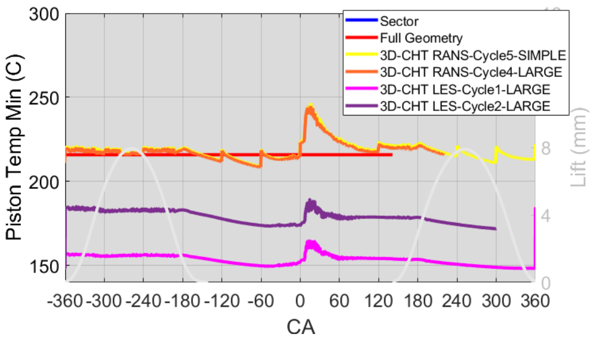

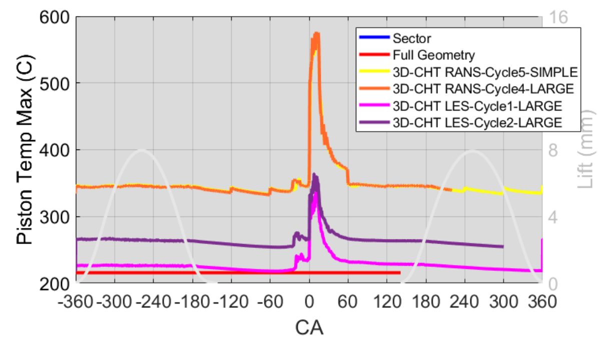

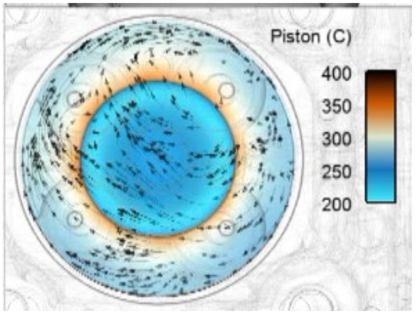

Head

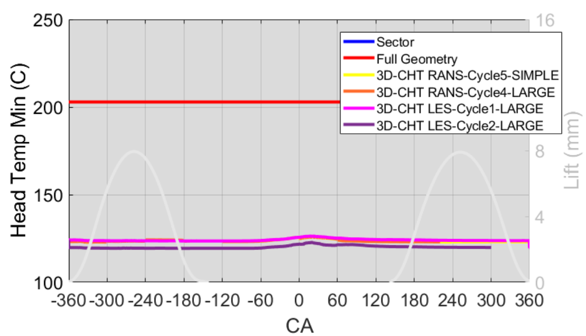

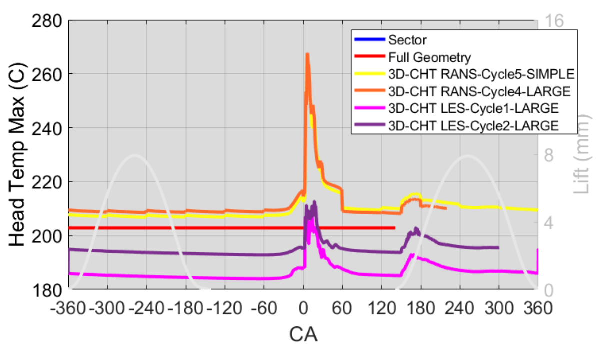

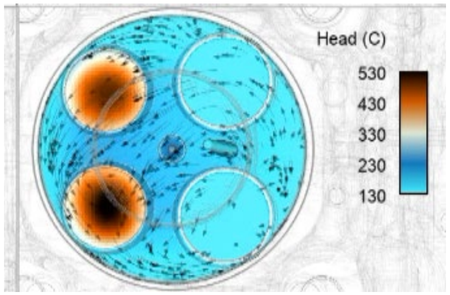

NOTE: Results shown are from a single case (2500 RPM, 19 bar)

CHT predicts significantly different thermal BCs… does that result in better emissions predictions?

Technical Accomplishments

Initial results: Emissions comparisons

	Baseline	Best Sector	Full-cyl	CHT+RANS
Geometry	Sector	Sector	Cylinder	Engine
Cycle	Closed	Closed	Open	Open
Max # cells	~170k	~1M	~2.5M	~4.5M
# species	47	144	144	144
# reactions	74	900	900	900
NOx	Zel'dovich	GRI 3.0	GRI 3.0	GRI 3.0
Soot	Detailed PSM	Hiroyasu	Hiroyasu	Hiroyasu
Turbulence	RANS	RANS	RANS	RANS
Other changes		Spray and wall-film	Intake swirl vane	CHT
Wall time / cycle	~2 hr	~5 hr	~3.5 days	~2 weeks
Cases	500	602	20	8

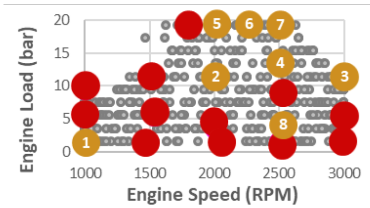

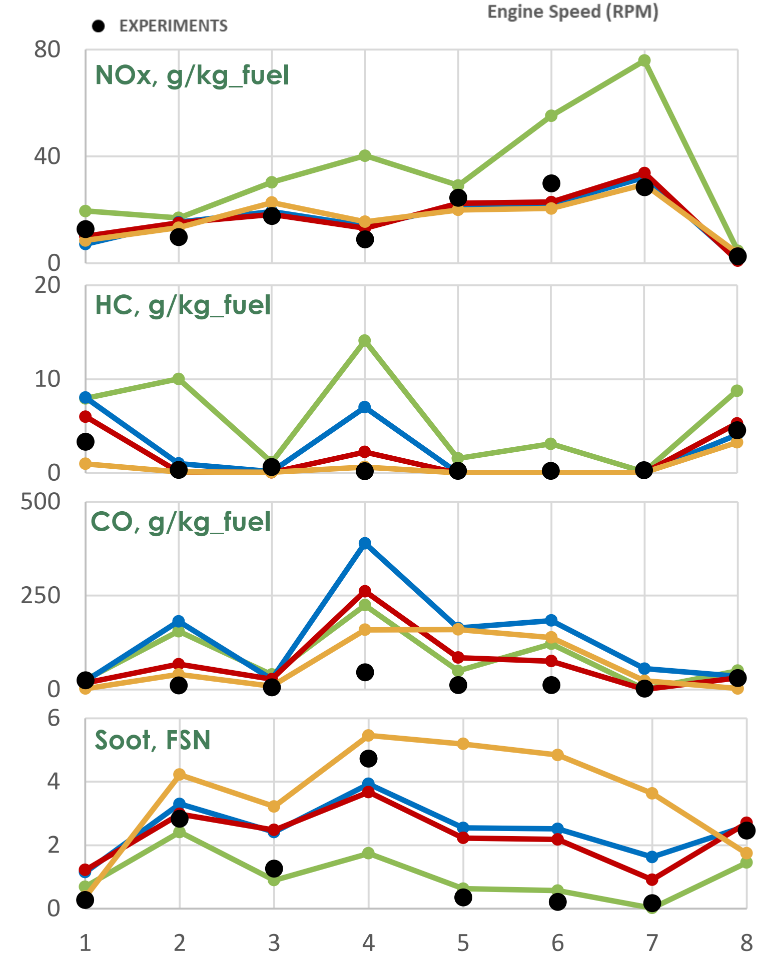

- Limited change in emissions with CHT+RANS despite changes in thermal BCs
  - PCP and CA50 (not shown) and NOx essentially same
  - HC slightly improved at for some cases
  - CO same or slightly improved for most cases but worse for others
  - Soot worse for most cases

Further analysis underway to fully understand impacts, CHT+LES simulations are still in progress

Milestones

Date	Milestone	Status
FY2018 – Q4	Evaluate impact of enabling increased detail in CFD engine simulations with HPC on predictive accuracy and computational requirements	**Met**
FY2019 – Q1	Complete refined CHT simulations and assess impact on model accuracy and computational requirements	**Met**
FY2019 – Q2	Develop LES model for the light-duty engine	**Met**
FY2019 – Q3	Perform initial runs of the light-duty LES engine model at a limited set of operating conditions	On Track
FY2019 – Q4	Assess impact of CHT and LES refinements on overall simulation accuracy and computational requirements	On Track

Future work

2018 ALCC project (GM, ORNL, LLNL, CSI) on virtual calibration continues until end of June

- Continuing to run CHT+LES simulations to generate as many cycles as possible
- Post-processing and data analysis continue

Expanding efforts to use CHT to examine engine knock

- Evaluate impact of spatio-temporal thermal BCs predicted by CHT to improve prediction of localized hot-spots and engine knock
- OLCF project with FCA and CSI has been on hold during additional model development
  - Currently ramping back up this spring
- Proposed new 2019 ALCC project (GM, ORNL, LLNL, CSI) to use best-practices approach developed with virtual calibration study to study engine knock

Transition to Summit

- Summit came online June 2018, opened for projects in January 2019
  - 200 PF, currently #1 on Top500 Super Computer list
  - Titan will be decommissioned at end of June 2019
- Working with CSI and LLNL to adapt CONVERGE v3.0 and Zero-RK for new architecture
  - Evaluate performance of enhanced HPC capabilities of CONVERGE v3.0

New efforts proposed under Combustion Consortium

- Collaborating with experimental efforts (ORNL, SNL) to develop validation methods for CHT
- Improve emissions prediction for cold-start and engine transients
  - Collaborating with experimental and diagnostic efforts (ORNL)

* Proposed future work is subject to change based on availability of funding and allocation of HPC resources

Response to reviewer comments

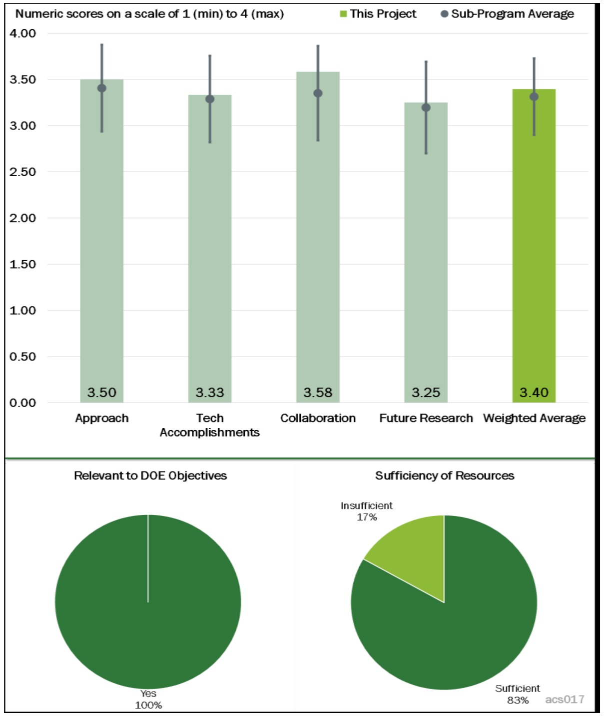

Many reviewer comments stressed importance of DOE support for collaboration and use of HPC to address issues facing industry

- "An outstanding approach only possible through HPC capabilities at national laboratories."
- *"The project is addressing problems relevant to OEMs."*
- "The approach of working with collaborators to push the limits of HPC by evaluating the computational performance of ever more detailed modeling efforts on realistic industry challenges is excellent. "
- "The team has pursued and shown success in a broad range of realistic problems such as knock modeling, viral design, LES, and conjugate heat transfer (CHT)."
- "This project shows to industry what is possible by applying HPC to practical problems that industry is facing."
- *"This is the type of work that national laboratories should be doing."*

Reviewer commented on need for more HDD-focused work

- **RESPONSE:** Efforts on an HD dual-fuel application under this project recently concluded (2017). We welcome additional collaborations with HD OEMs in the future. HD modeling efforts at ORNL supporting the Daimler SuperTruck II team will be presented under **ACE100**, and MD/HD modeling efforts supporting Co-Optima will be presented under **FT071**.

Reviewer suggested using this approach (including CHT) to investigate cold-start and transients

- **RESPONSE:** We agree and, as discussed in future work, have begun efforts to study cold-start emissions as a sub-task under this project in collaboration with experimental efforts at ORNL under the combustion consortium.

Collaborations and Acknowledgments

Project supports collaborative efforts with industry, NLs, and universities to apply DOE ASCR HPC resources to accelerate development of advanced engines capable of meeting fuel economy and emissions goals

- DOE Office of Energy Efficiency and Renewable Energy
  - Vehicle Technologies Office – Gurpreet Singh, Michael Weismiller
- DOE Office of Science
  - Advanced Scientific Computing Research program
  - OLCF User Facility – Jack Wells, Ashley Barker, et al.
  - ACCEL Industrial Partnership Program – Suzy Tichenor
- LLNL – Russell Whitesides, Matt McNenly
- GM – Ronald Grover, Jian Gao, Venkatesh Gopalakrishnan, Ramachandra Diwakar
- CSI – Dan Lee, Keith Richards, Jon Povich, Tristan Burton, Nitesh Attal, Sameera Wijeyakulasuriya, et al.
- FCA – Gary Cai, Ken Hardman, Ron Reese

Enabling virtual engine design and calibration with HPC and GPUs	
General Motors	Provide technical direction, engine model, validation data
Convergent Science	Support CONVERGE on OLCF machines; assist with model development
Oak Ridge National Laboratory	Support simulation efforts and HPC resources; provide analysis and direction
Lawrence Livermore National Laboratory	Support use of Zero-RK GPU-enabled chemistry solvers
Investigating initiation of autoignition and knock in GDI engines	
Fiat Chrysler Automobiles	Provide technical direction, engine model, validation data
Convergent Science	Support CONVERGE on OLCF machines; assist with model development
Oak Ridge National Laboratory	Support simulation efforts and HPC resources; provide analysis and direction

Remaining challenges and barriers

Accurate prediction of soot remains a challenge

- Use of the detailed particle-size mimic (PSM) soot model in CONVERGE has provided better results at some operating points
  - Exists as part of SAGE chemistry solvers, cannot be used with Zero-RK UDF
- Use of additional phenomenological soot models as UDFs may be required

CHT model produces different thermal BCs, but have no direct method to validate accuracy

- Proposed efforts under the Combustion Consortium plan to develop approaches for gathering validation data for thermal BCs

Applying simulation tools on HPC resources presents unique challenges

- Scalability is vital to maximize benefits of parallel architectures
  - Working with CSI to evaluate enhanced HPC capabilities of CONVERGE 3.0 on Summit
- Increased model fidelity results in increased memory demands
  - Summit has ~600GB/node (compared to 32 GB/core on Titan) which should help with this
- Maximizing benefits of HPC (increased detail, reduced wall time, etc.) must often be balanced with administration rules (queue and scheduling rules, fixed allocations, I/O limitations, etc.)
- Software must continually adapt to evolving hardware technologies
  - Codes need to be increasingly hardware independent
  - OLCF projects must make use of Titan's GPUs
  - Forthcoming DOE-ASCR machines (Summit, Aurora) will further exacerbate this issue
- *Partnering with ISV is crucial for addressing many of these issues*

Summary

Relevance

- Supports collaborative efforts using DOE ASCR HPC resources to accelerate development of advanced engines with improved efficiency and emissions

Approach

- HPC and GPU-enabled solvers enable higher fidelity models and faster throughput
- Systematic approach to add model detail and assess impacts on accuracy and speed

Technical Accomplishments

- Adding detail to the CFD engine models has provided significant improvement in accuracy over the full operating range without tuning for individual cases
  - Remaining issues include the more extreme calibration points at high speed and load and soot
- Initial results show significant differences in initial and boundary conditions with full-cycle model and CHT than those obtained from calibration with lower-order models
- Initial results suggest that CHT+LES predicts lower thermal diffusion rates than CHT+RANS resulting in slower convergence and smaller wall temperature swings (real or numeric?)

Collaborations

- Collaborative efforts with OEMs, ISVs, and NLs
- Leverage of DOE's ASCR leadership computing resources at OLCF User Facility

Future Work (subject to change based on availability of funding and allocation of HPC resources)

- Expanded efforts using CHT+LES to examine engine knock
- Develop validation approaches for CHT predictions
- Develop improved approaches for predicting emissions during cold-start and transients

Backup Slides

"Supercycling"

- Combustion model converges to steady-state solution much faster than thermal model
  - MANY combustion cycles would need to be simulated to reach thermal steady state
- CONVERGE uses supercycling to resolve this issue
  - Combustion and CHT are solved at every time step
  - Gas-side thermal boundary data (near-wall T, HTC) are cycle-averaged over moving, 720-degree window
  - Beginning at the end of the first cycle and repeated at user-specified intervals…
    - Cycle-averaged BCs are applied to the solid model
    - Solid model is run to steady state
- Usually requires ~3 cycles to converge