Turbulent Combustion Simulation and *In Situ* Analytics on Titan and Summit with S3D-Legion

Jacqueline H. Chen[1], Elliott Slaughter[2], Wonchan Lee[3], Sean Treichler[4], Michael Bauer[4], Konduri Aditya[1], Hemanth Kolla[1], Pat McCormick[5] and Alex Aiken[3]

[1]Sandia National Laboratories

[2]SLAC National Accelerator Laboratory

[3]Stanford University

[4]NVIDIA

[5]Los Alamos National Laboratory

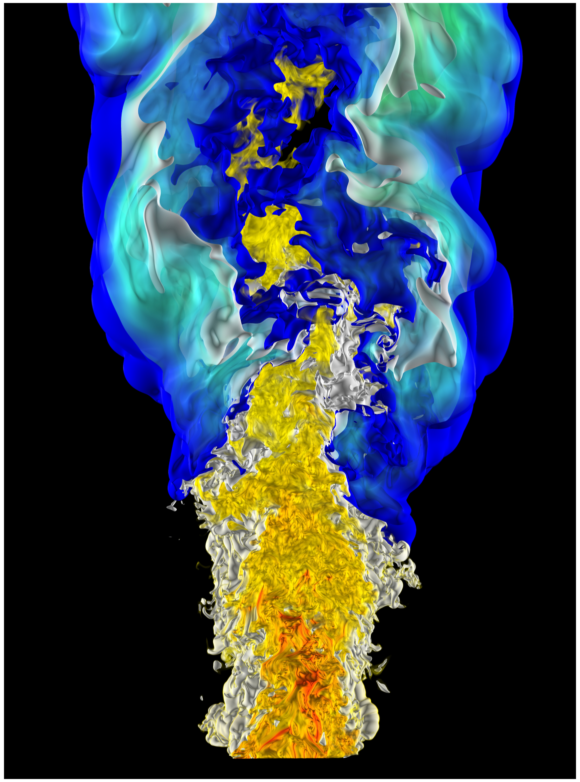

1st R-CCS Symposium, February 18-19, 2019

Office of Science
U.S. Department of Energy

Sandia National Laboratories

Office of Science
INCITE
Advancing America's Science and Industrial Competitiveness

ECP
EXASCALE COMPUTING PROJECT

Exascale Simulation of Combustion

- Predict behavior of new fuels in different combustion scenarios at realistic pressure and turbulence conditions
  - Optimize co-design of fuels and engines
  - Reduce greenhouse gas emissions, increase efficiency, shorten time-to-market
- High-fidelity direct numerical simulation and hybrid DNS/LES methodologies
  - turbulence- chemistry interactions
  - chemical fidelity to differentiate effects of fuels
  - uncertainties in thermo-chemical properties and physics models
  - complex flows due to compression by a piston, swirl, bluff-bodies, cavities

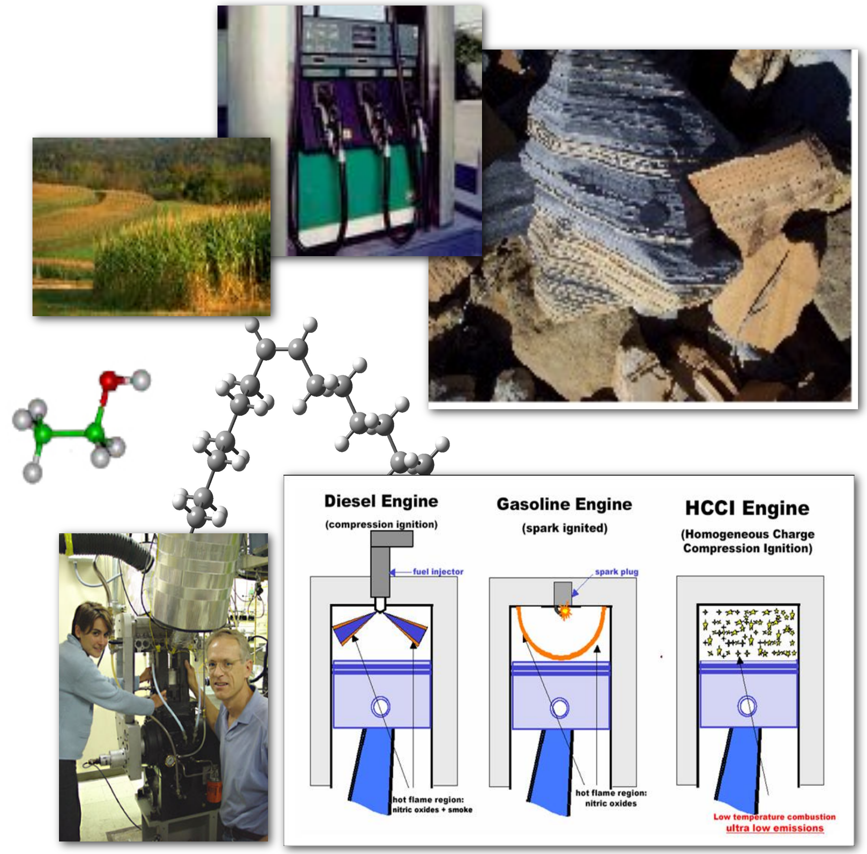

Emissions and Efficiency Drivers for Autoignition in Engines

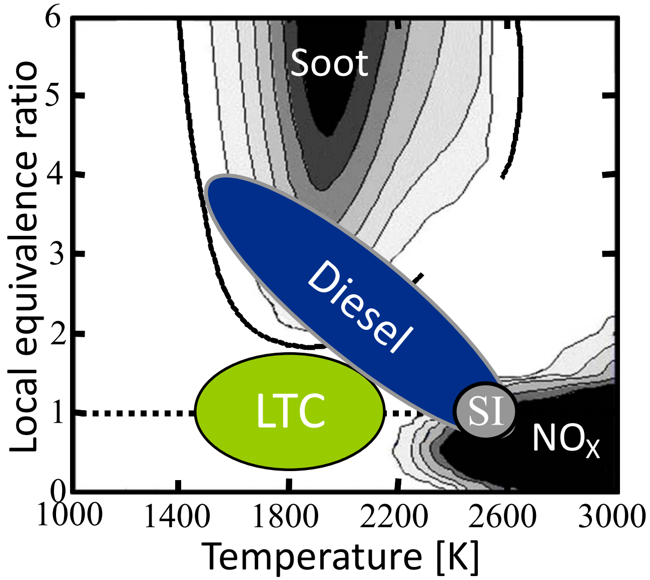

Conventional diesel *

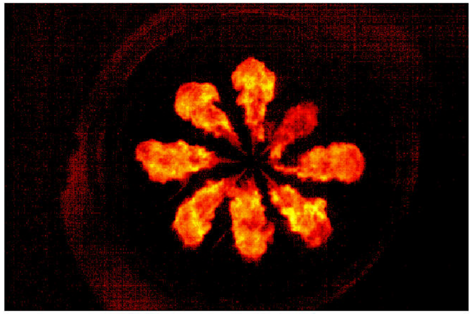

Early injection PCCI

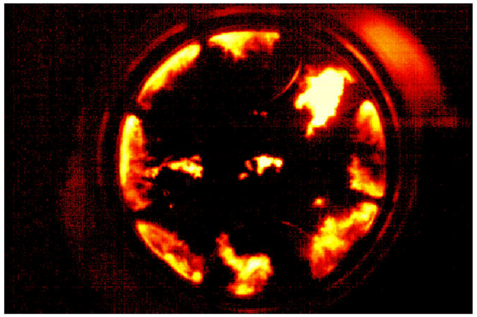

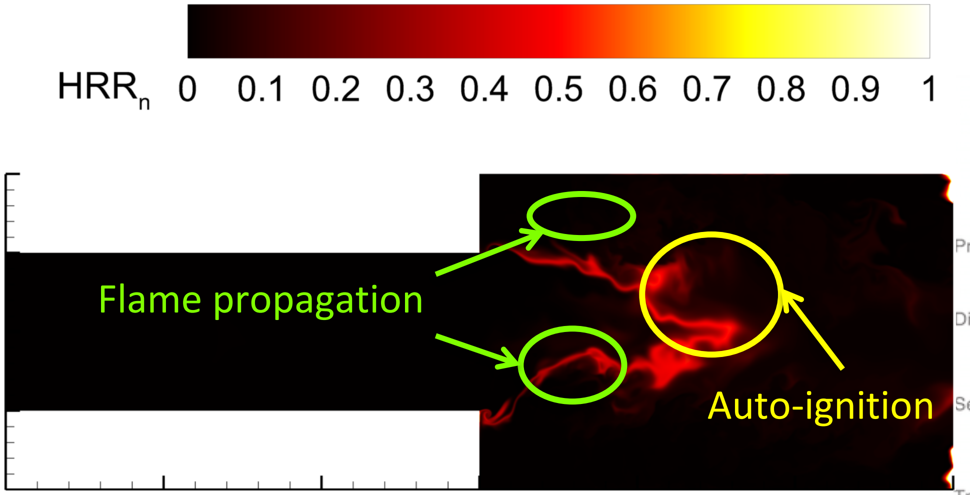

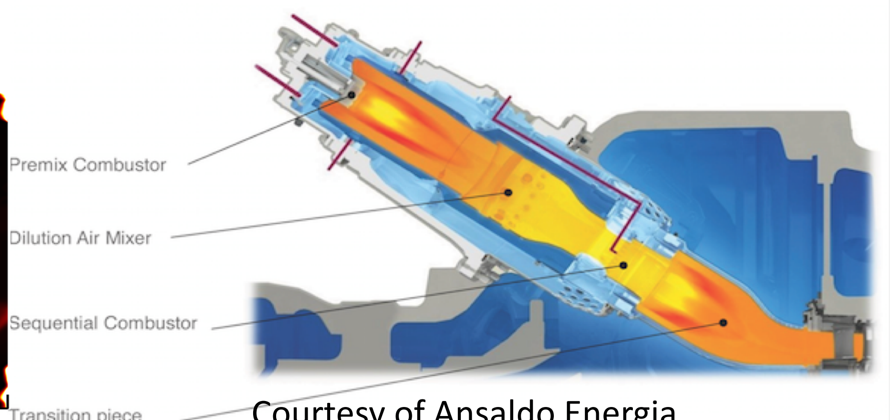

Courtesy of Ansaldo Energia

Fundamental Turbulence-Chemistry Interactions Motivated by Advanced Engines

- Higher fuel efficiency and lower emissions driving combustion towards more **fuel lean, partially-premixed conditions**

- **Mixed-regime combustion**

- **Multi-stage autoignition - sensitivity to fuel chemistry and preferential diffusion**

- **Multi-scale interactions and backscatter**

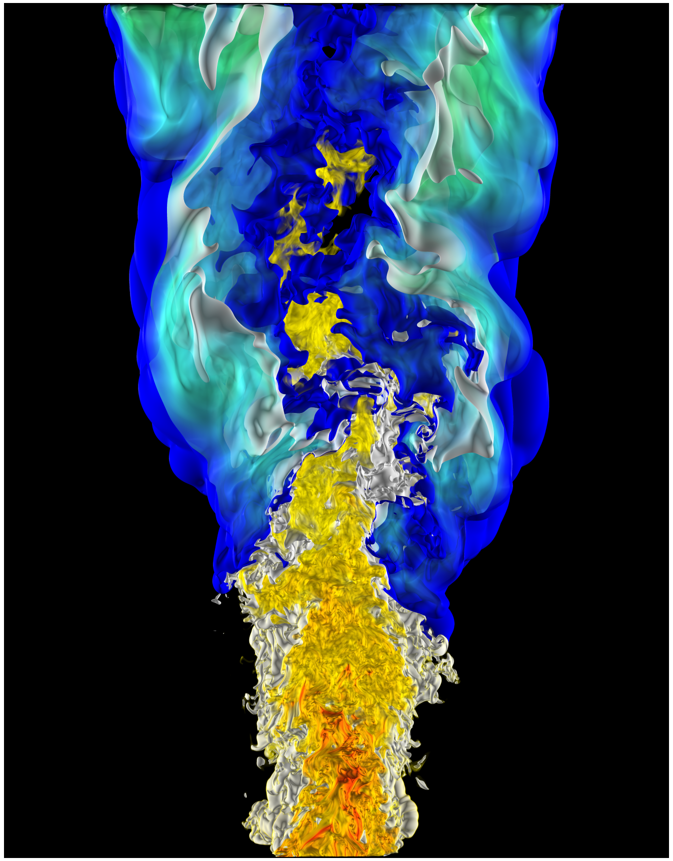

Lifted dimethyl ether jet flame at 5 bar stabilized on cool flame ignition intermediate, flame marker OH (blue); methoxymethylperoxy (ignition marker (yellow)

DOE Exascale Computing Project (ECP) uses co-design and integration to achieve capable exascale in 2021-2023

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable and productive software stack	Hardware technology elements	Integrated exascale supercomputers

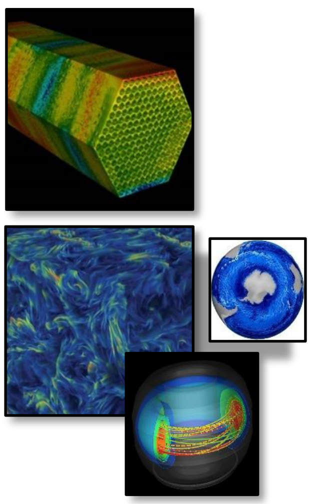

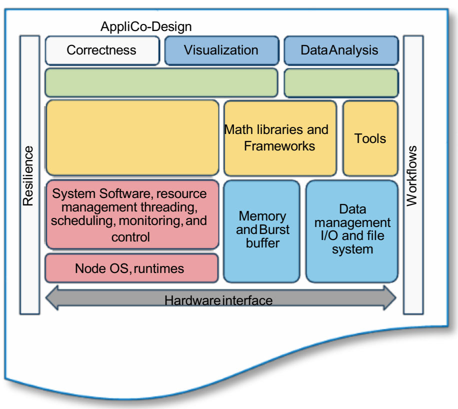

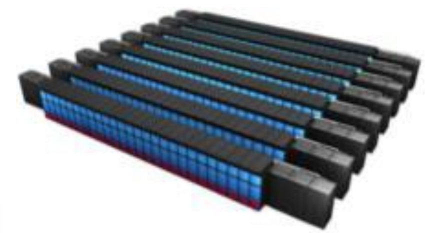

ECP's work encompasses applications, system software, hardware technologies and architectures, and workforce development

 From Paul Messina's ASCAC talk April 19, 2017

EXASCALE COMPUTING PROJECT

Pele Project: Transforming Combustion Science and Technology Through Exascale Simulation

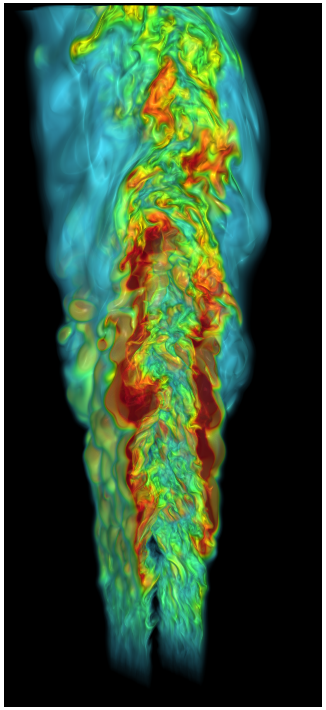

Effects of reactivity stratification at:

- high pressure
- high turbulence
- fuel blends

on:

- ignition delay
- combustion rates
- emissions

Pele: Block-structured adaptive mesh refinement, multi-physics: spray, soot, and radiation, real gas, complex geometry

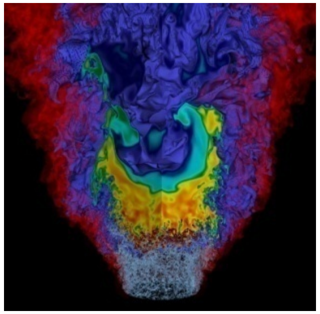

Automated Mechanism Generation

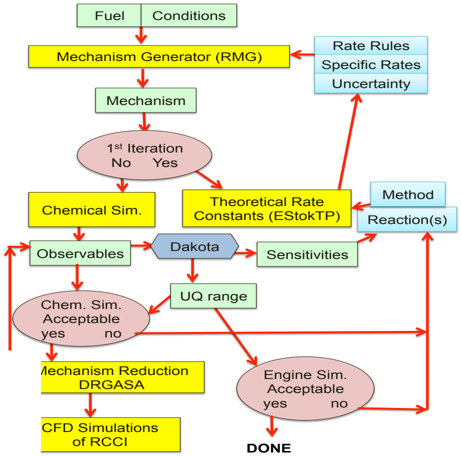

S3D: multi-block compressible reacting DNS multi-physics validation: spray, soot, radiation

Direct Numerical Simulation – S3D

- DNS of turbulent reacting flows
- Solves compressible reacting Navier-Stokes, total energy and species continuity equations
- High-order finite-difference methods
- Detailed reaction kinetics and molecular transport models
- Lagrangian particle tracking (tracers, spray, soot)
- In situ analytics and visualization
- Machine learning
- Refactored for multi-threaded, many core heterogeneous architectures

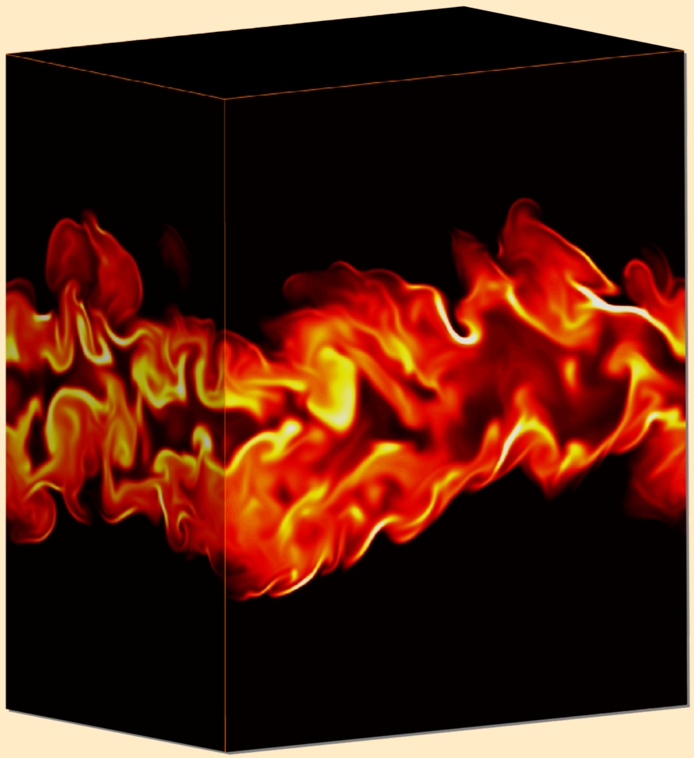

DNS provides unique fundamental insight into the chemistry-turbulence interaction

Chen *et al.*, *Comp. Sci. Disc.*, 2009, *Treichler et al. 2017*

Computational intensity of DNS scales with Moore's Law

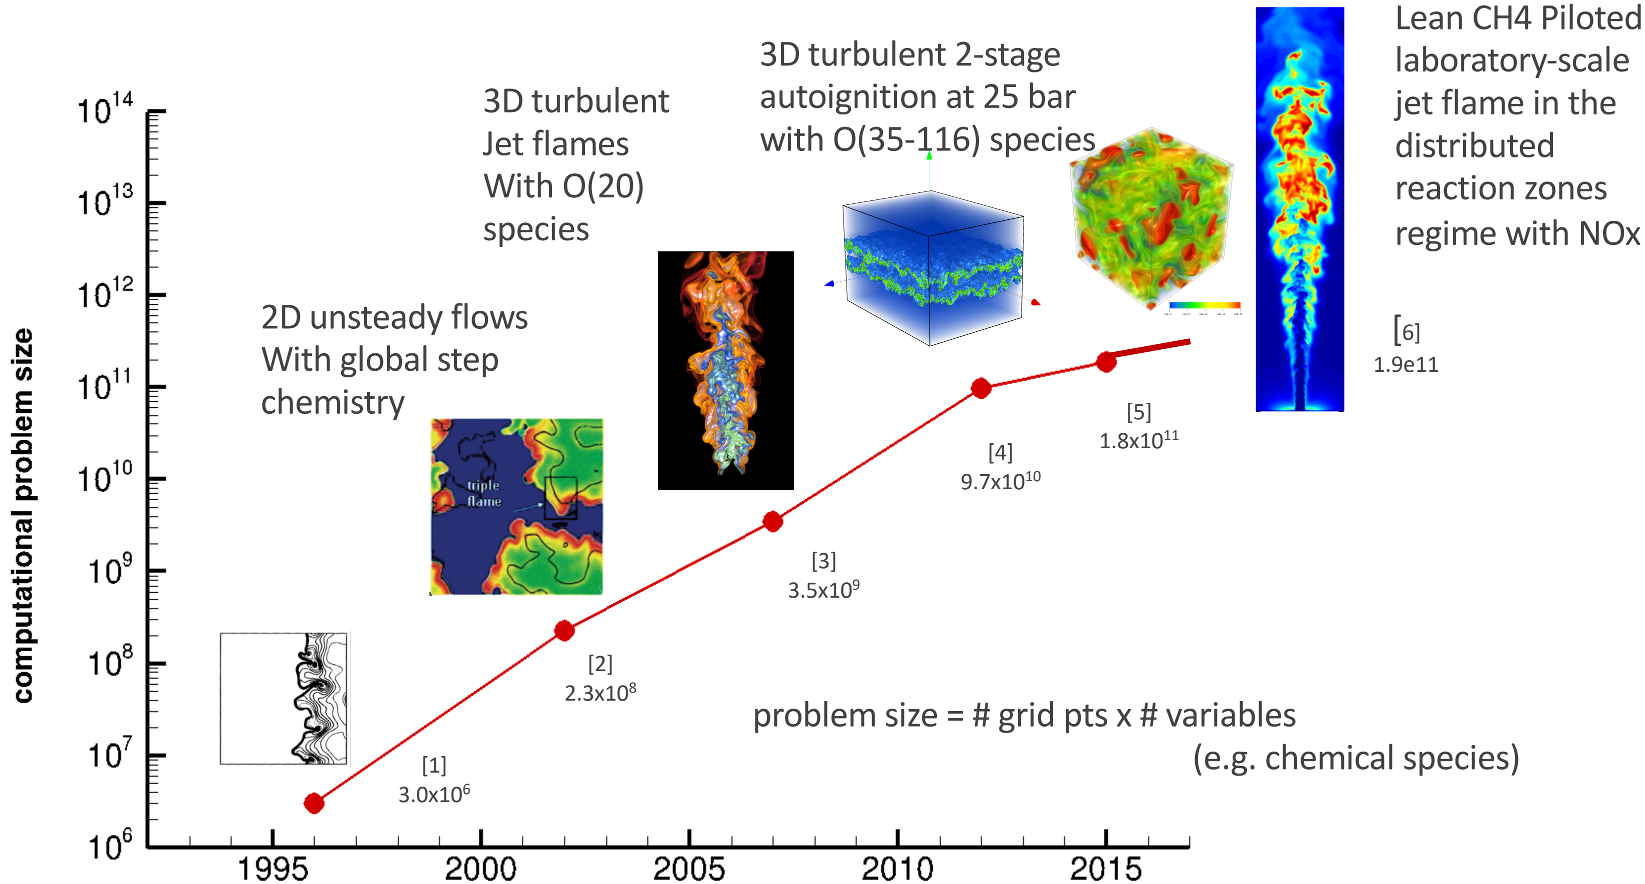

[1] T. Echekki, J.H. Chen, *Comb. Flame*, 1996, vol.106.
[2] T. Echekki, J.H, Chen, *Proc. Comb. Inst.*, 2002, vol. 29.
[3] R. Sankaran, E.R. Hawkes, J.H. Chen, *Proc. Comb. Inst.*, 2007, vol. 31.
[4] E.R. Hawkes, O. Chatakonda, H. Kolla, A.R. Kerstein, J.H. Chen, *Comb. Flame*, 2012, vol. 159.
[5] 2015 submission for Gordon Bell prize
[6] H. Wang, E. Hawkes, J. H. Chen, Comb. Flame 2017

Outline:

- DNS of tubulent combustion with complex flow at the petascale:
  - DNS of turbulent autoigniting n-dodecane diesel jets
  - DNS of reheat combustion in staged gas turbines
- Path to exascale combustion simulations (2021-2023)
  - Programming models
  - Composable in situ workflow (analytics, machine learning)

DNS of a Turbulent Autoigniting n-Dodecane Jet at 25 Bar

Giulio Borghesi[1], Alex Krisman[1], Tianfeng Lu[2] and Jackie Chen[1]

[1]Combustion Research Facility, Sandia National Laboratories

[2]University of Connecticut

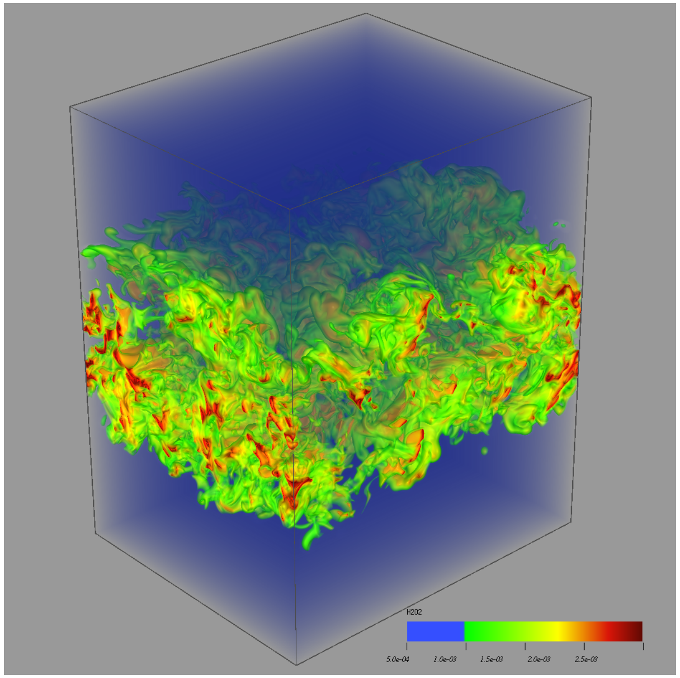

Ketohydroperoxide mass fraction

Borghesi et al. Combustion and Flame, 2018

Background and Objective

- Low-temperature combustion (LTC) aims at increasing fuel efficiency and reducing emissions
- Under LTC conditions, combustion occurs in a mixed mode and in multiple ignition stages
- Ignition is now very sensitive to the fuel chemistry, especially to the low temperature reactions branch

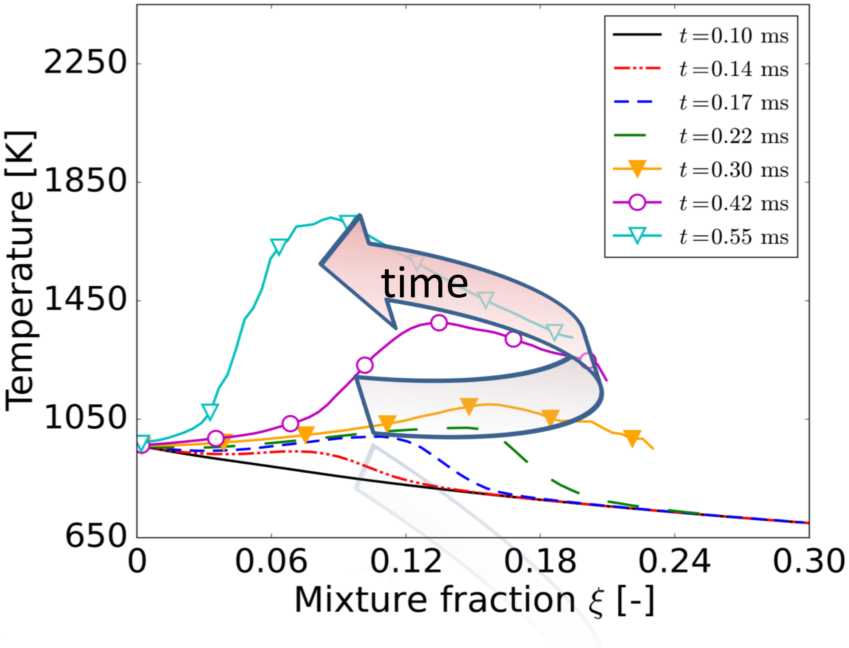

Question: How does transport and low-temperature chemistry affect ignition in low-temperature diesel combustion?

Background on high-temperature ignition

- Homogeneous PSR calculations show the existence of x value, x_m, where ignition delay has a minimum;
- Flamelet simulations show the ignition delay increases with scalar dissipation N until a critical value is reached;
- In practical systems, ignition occurs at locations close to x_m where N is low. The ignition delay is longer than in PSR;
- **Question:** which features of high-T ignition carry over to low-T ignition?

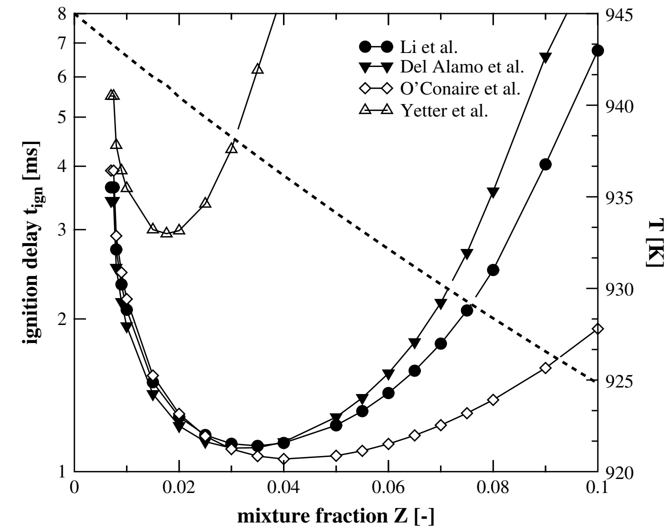

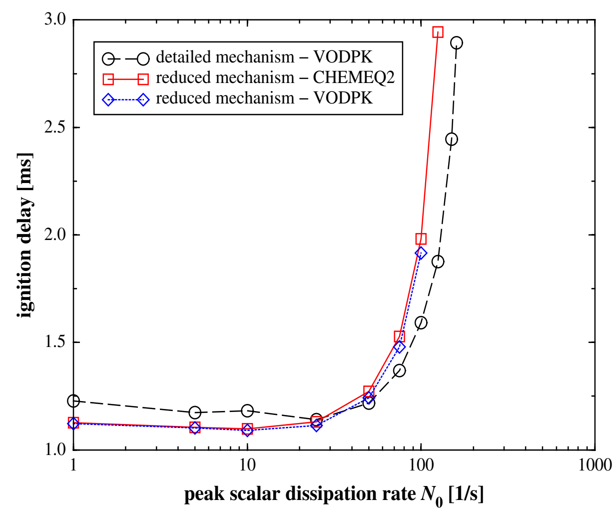

Figure: ignition delay time in PSR and in nonpremixed flamelet simulation[1]

[1] E. Mastorakos, PECS (2009), pp. 57-97

Low Temperature Diesel Combustion Experiments – Engine Combustion Network

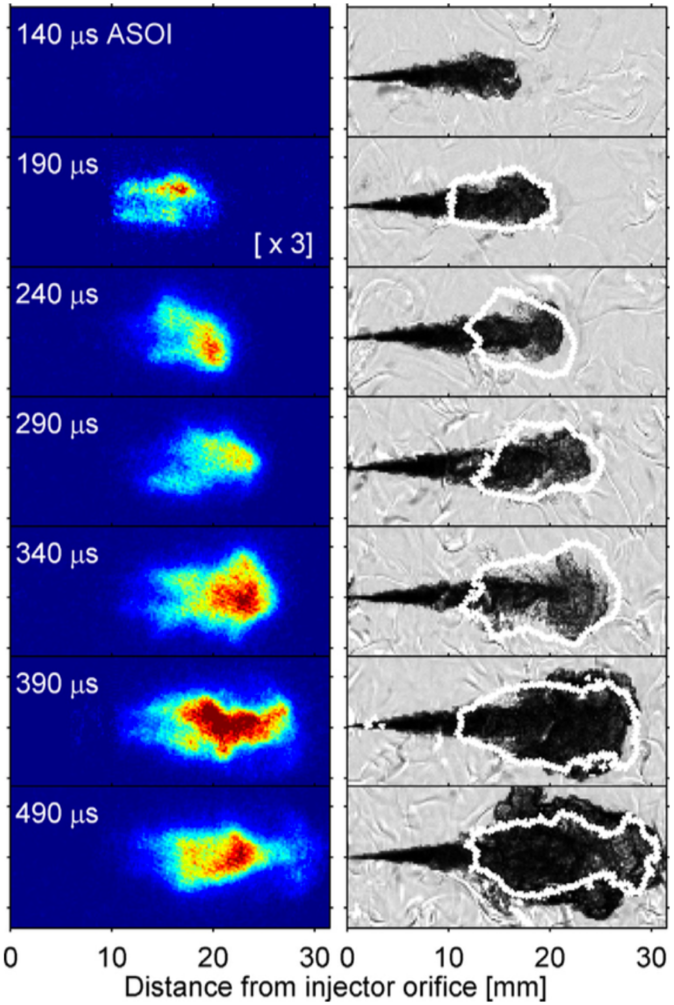

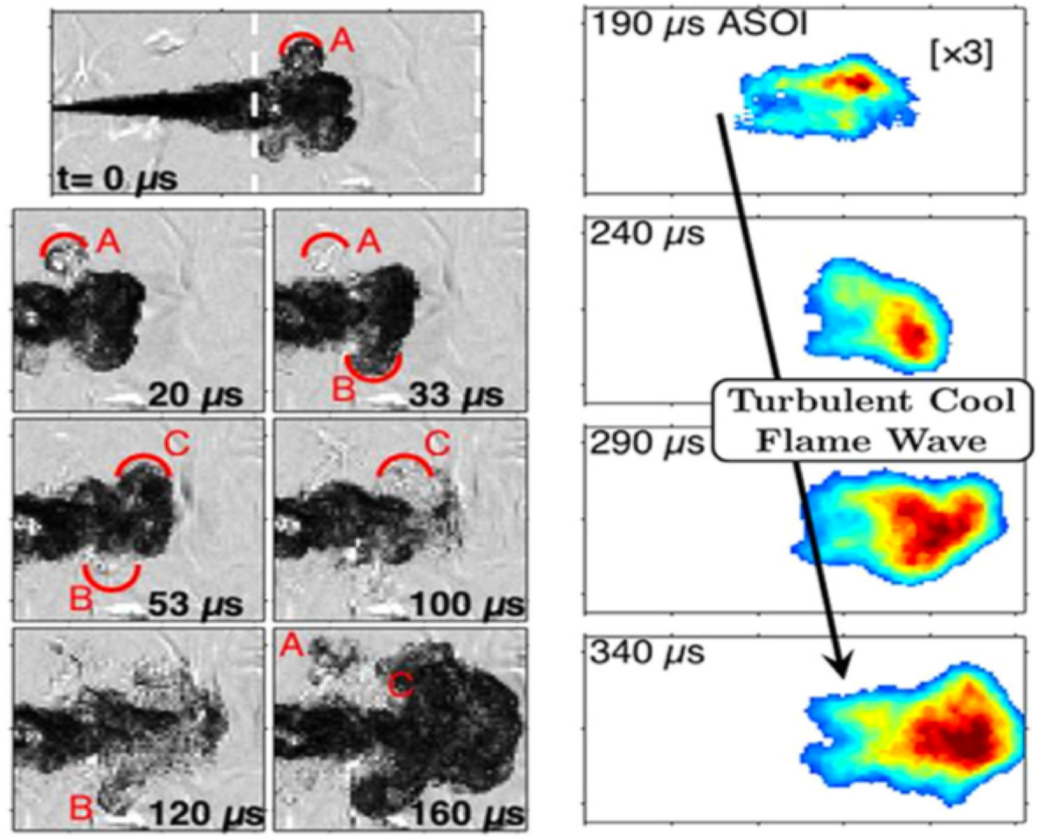

Skeen et al., PROCI 35 (2015) 3167-3174

Dahms et al., PROCI 36 (2017) 2615-2623

DNS Configuration and Physical Parameters

- **Pressure:** 25 bar
- **Air stream:** 15% X_{O2}+85% X_{N2}, T=960 K
- **Fuel stream:** *n*-dodecane at ξ=0.3, T=450 K
- **Kinetics:** 35-species non-stiff reduced (Lu)
- **Fuel jet velocity:** 21 m/s, Re_j = 7000, Re_t ~ 950
- **Code and cost:** S3D Legion, 60M CPUh
- **Setup:**
  - 3 billion grids
  - 3 microns spatial grid resolution
  - Dimensions: 3.6 mm x 14.0 mm x 3.0 mm
  - 1 ms of physical time with 4 ns timesteps to observe ignition and propagation of burning flames throughout the domain
  - BCs: X and Z periodic, Y NSCBC outflows

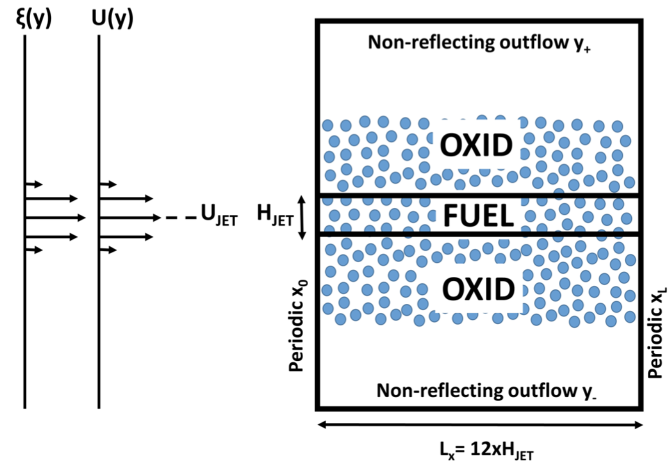

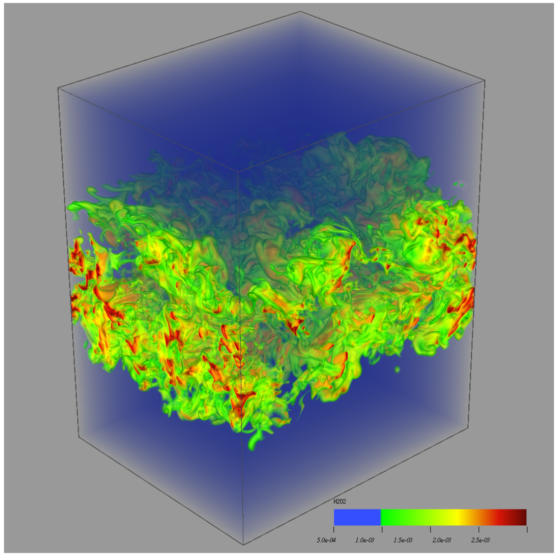

Figure: H_2O_2 mass fraction at t=0.17 ms after start of reactions

Homogeneous Multi-Stage Autoignition

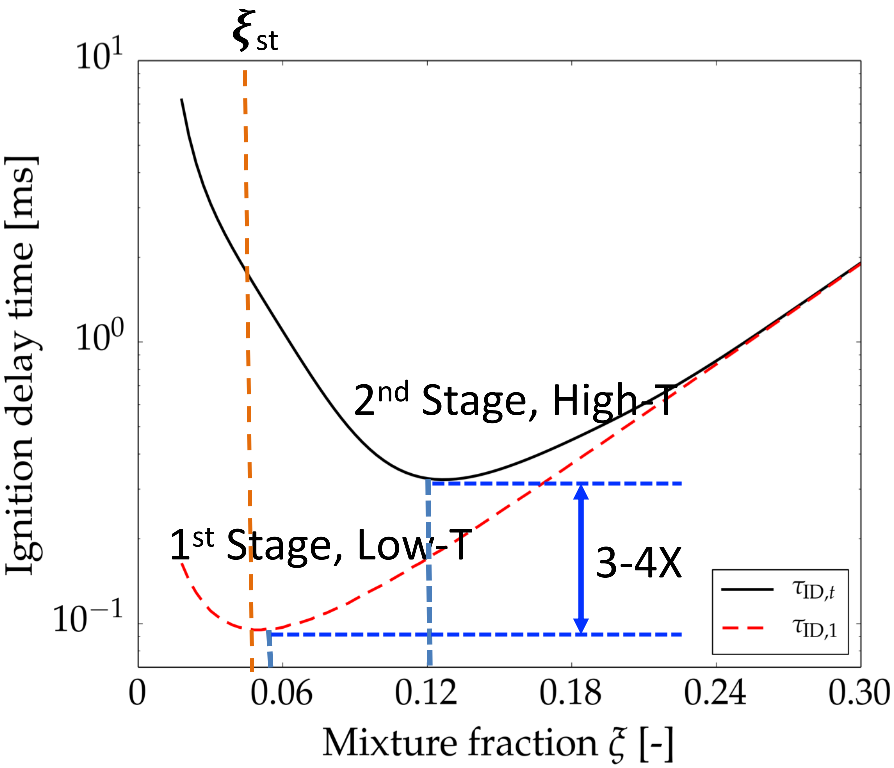

Homogeneous ignition delay time

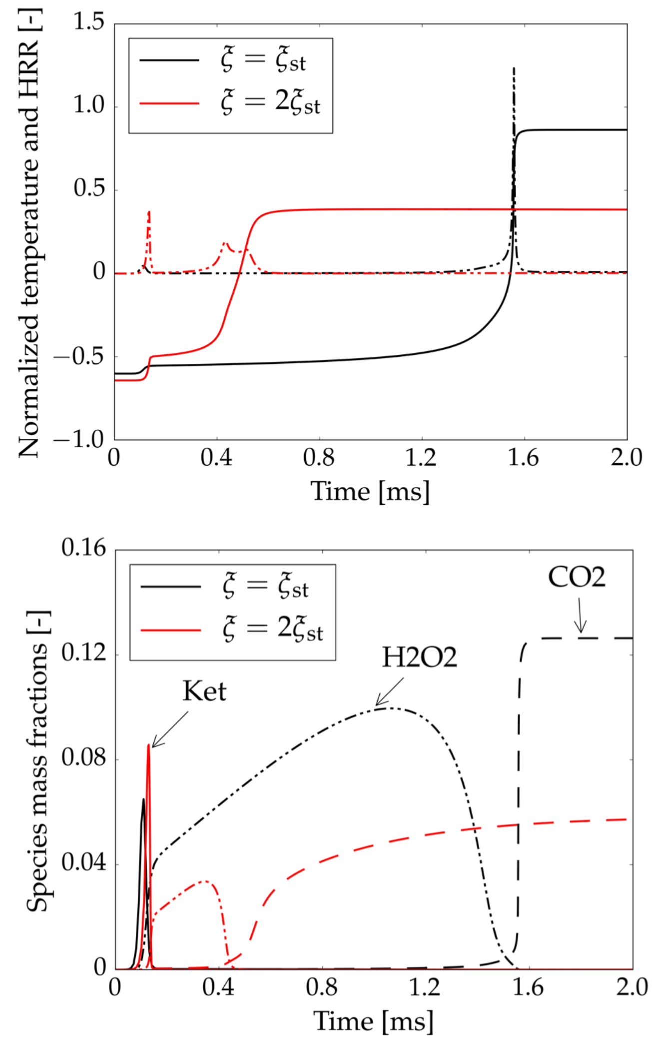

Temporal evolution of selected reactive scalars

Dynamics of Two-Stage n-dodecane Autognition in a Jet at Diesel Conditions

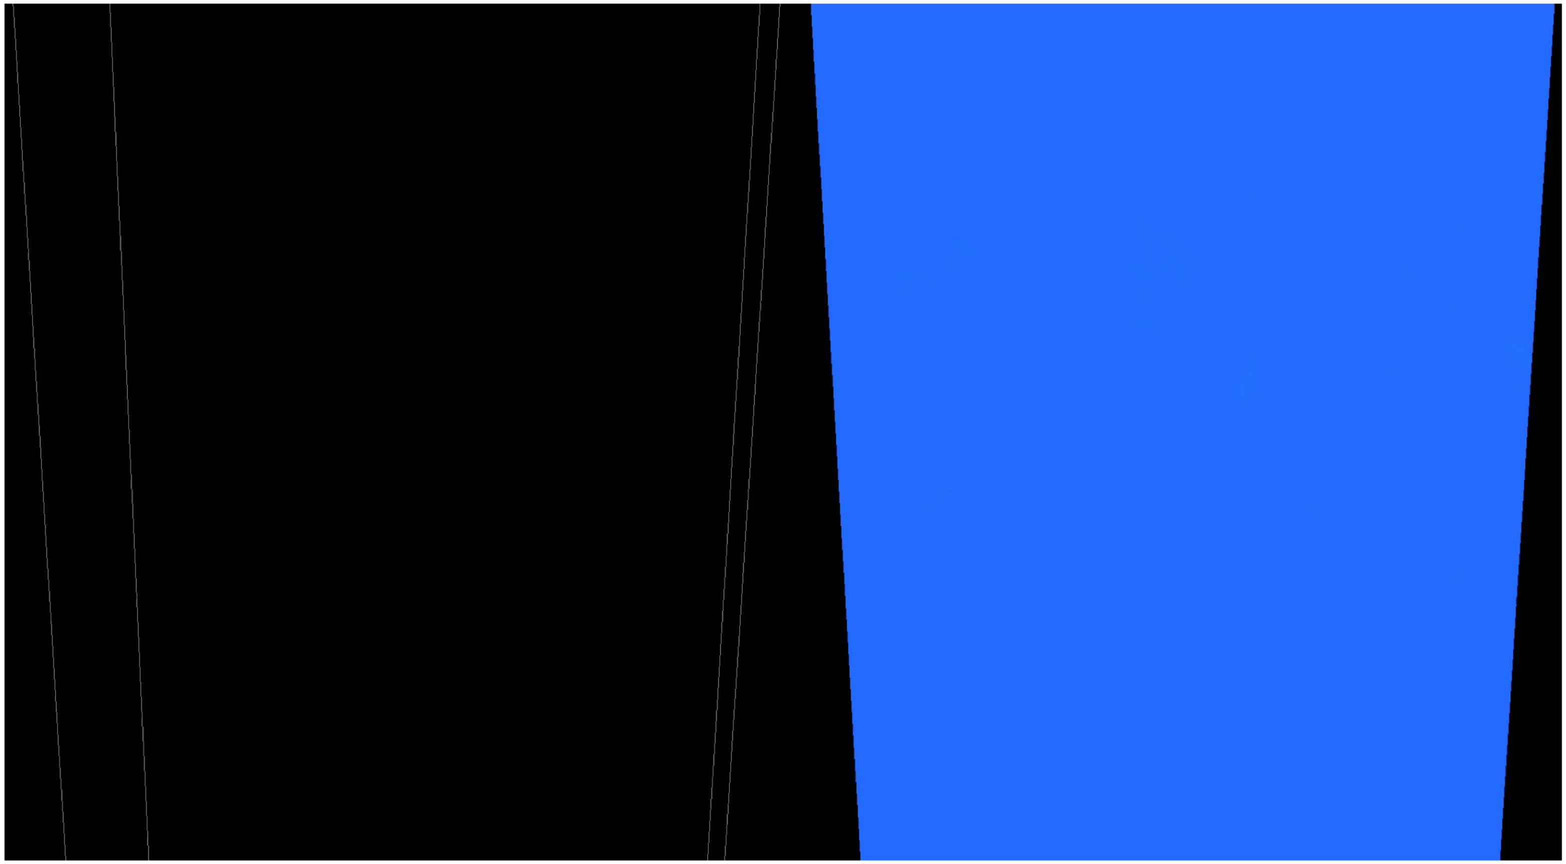

Rendering by Chris Ye, Min Shih, Franz Sauer, and Kwan-Liu Ma

Ketohydroperoxide and *T,K* (>1150K) **H_2O_2 mass fraction**

Conditional statistics reveal ignition dynamics

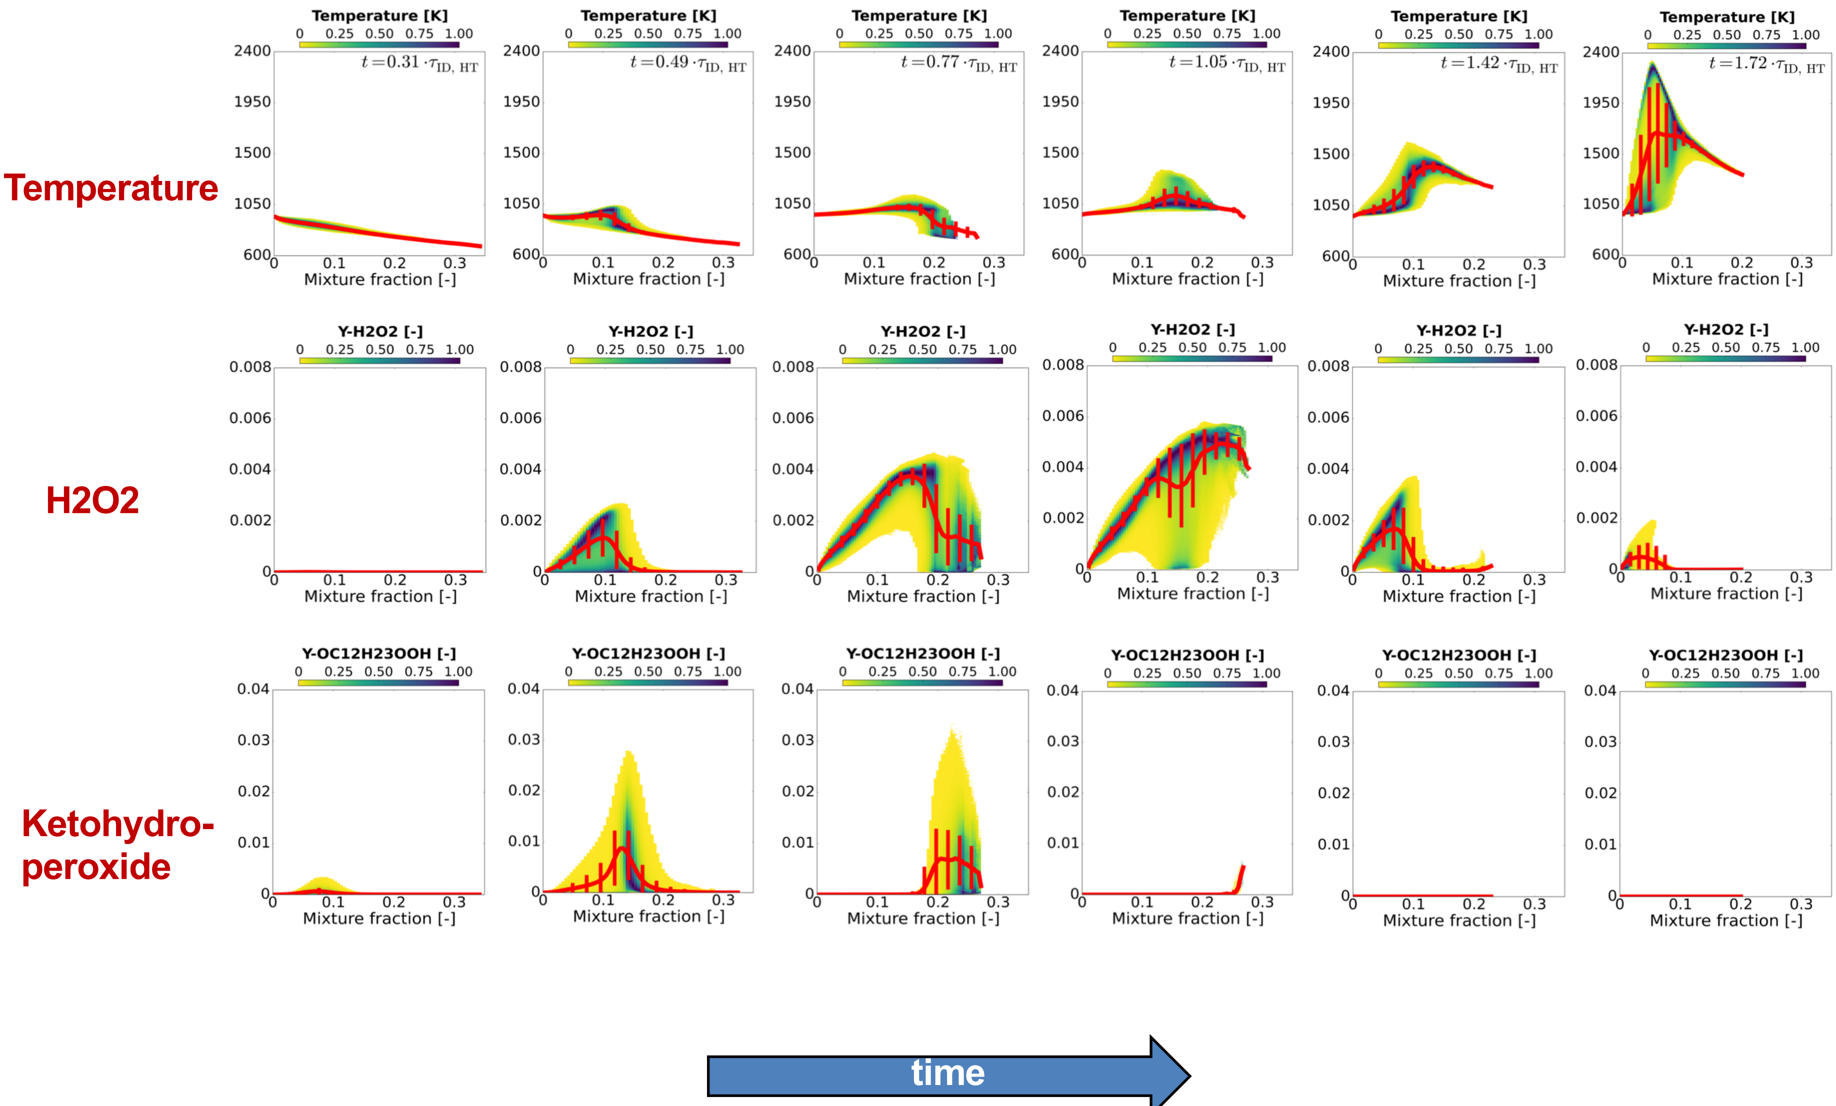

Turbulent versus homogeneous ignition

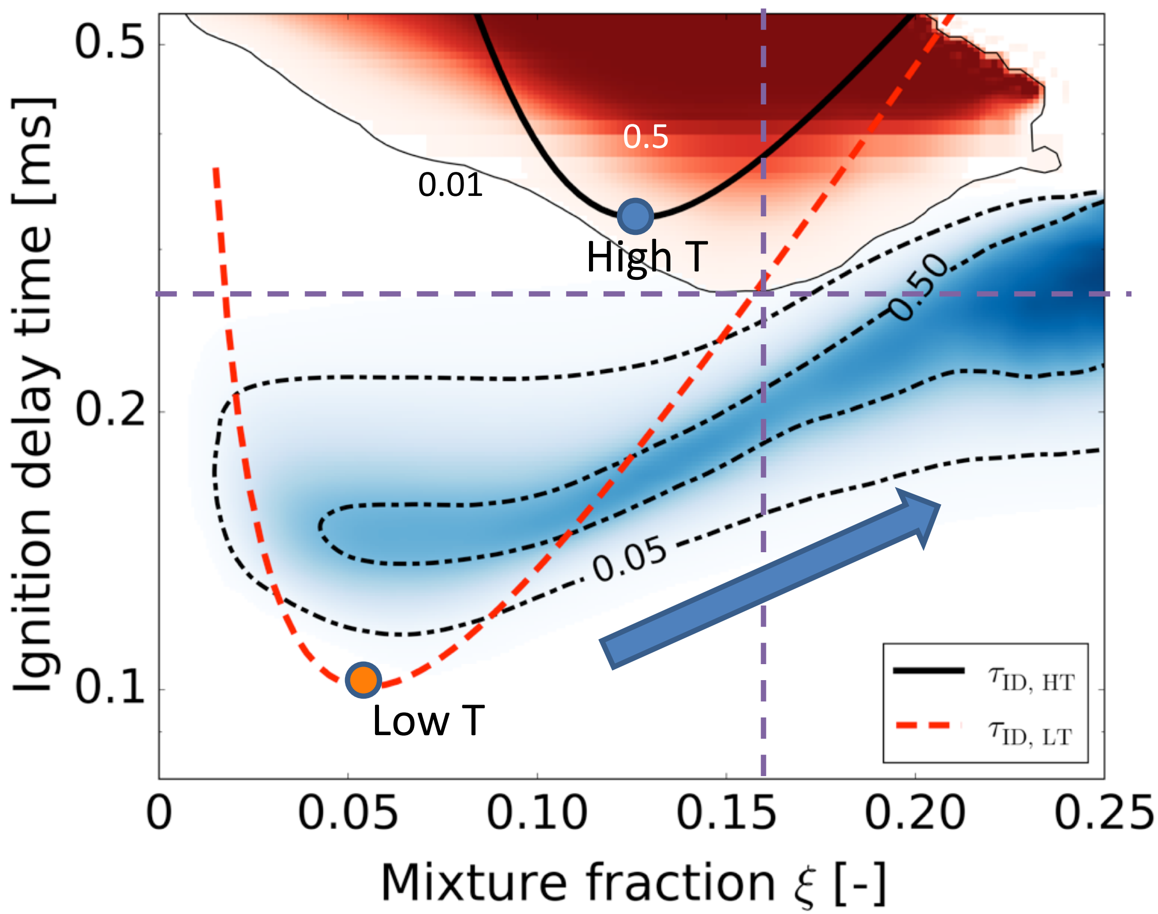

Low-T and high-T ignition in jet can be faster and than homogeneous ignition!

Propagation mechanism for low-T reaction front

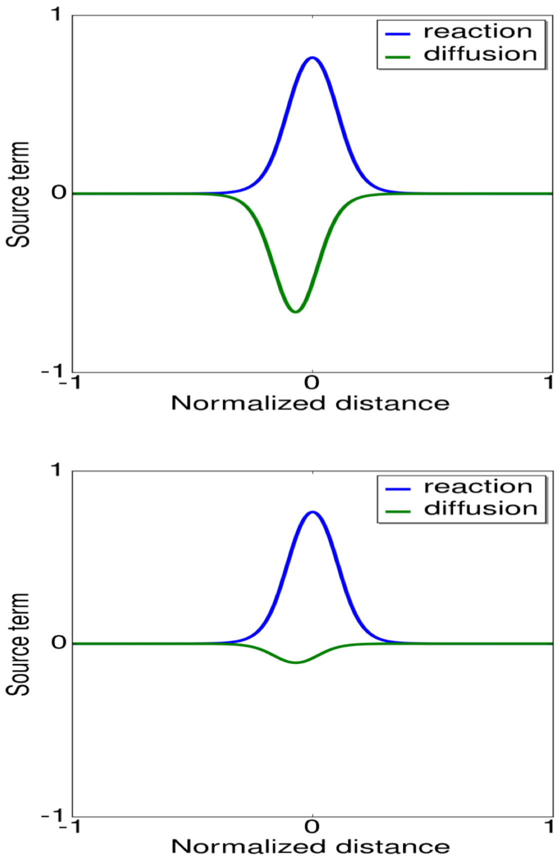

Sketch of reaction / diffusion balance along normal for KET for flame and ignition kernel

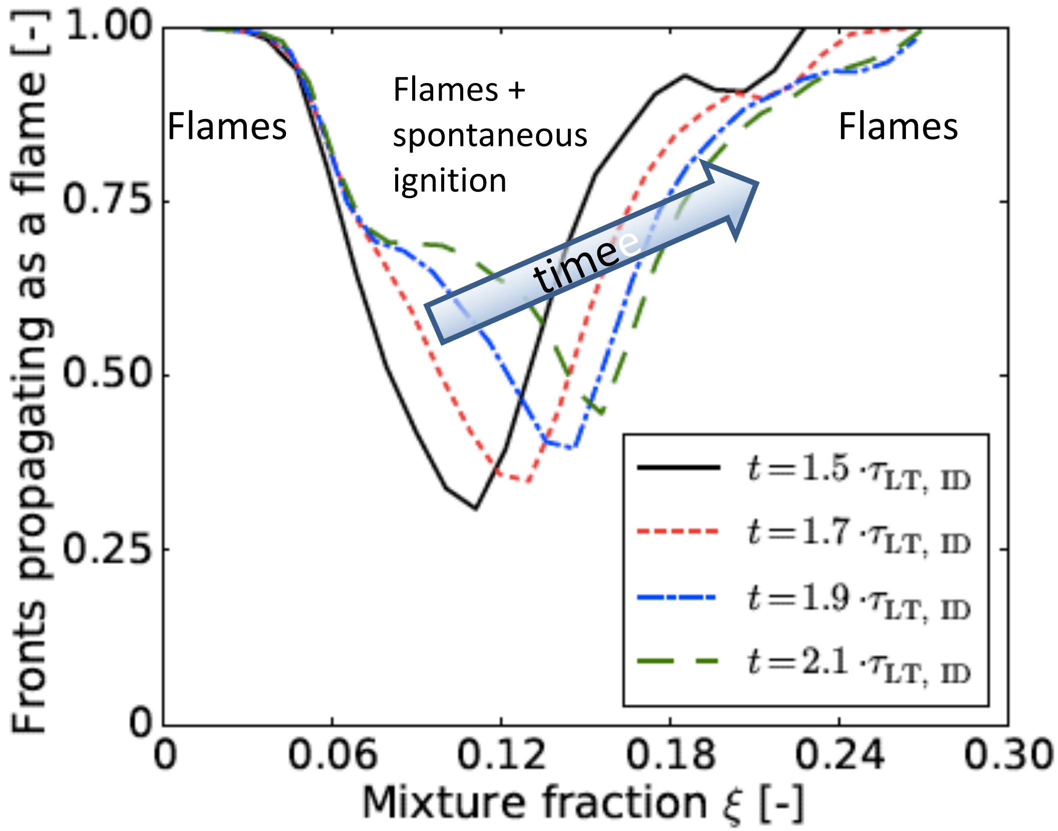

Figure: Fraction of reactive low-T fronts propagating as a flame

Low-T fronts propagate through diffusively supported flame

Effect of Scalar Dissipation Rate on Low Temperature Ignition

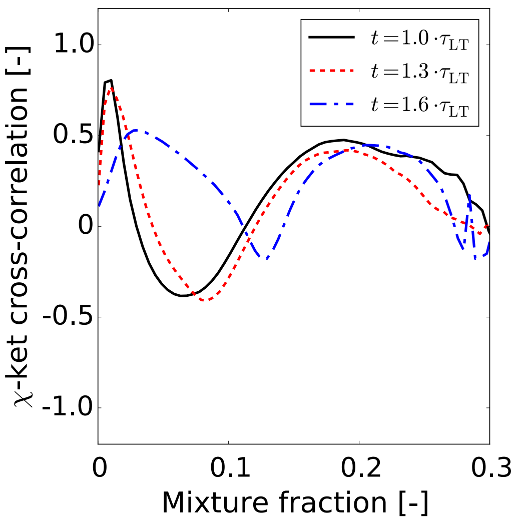

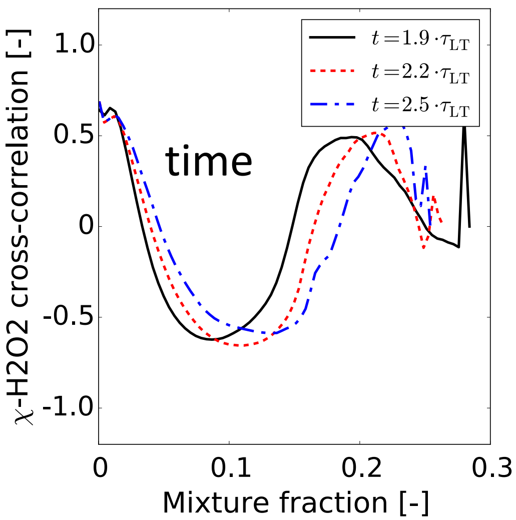

Conclusions

- Low-temperature reactions create the conditions for high-temperature ignition to occur faster than under homogeneous conditions;
- Low-temperature front appears to propagate through a diffusively supported cool flame;
- High scalar dissipation appears to delay low-temperature ignition; however, it leads to faster ignition at very rich mixture conditions;
- High-T ignition starts at conditions richer-than-homogeneous conditions (ξ=0.16 compared to ξ=12). Edge flames are seen to form around ξ_{st}. High-T flame ignites mainly by propagation of rich premixed flames following hot ignition to ξ_{st}.

Direct Numerical Simulation of flame stabilization assisted by auto-ignition at *reheat* conditions

Konduri Aditya[a], Andrea Gruber[b], Mirko Bothien[c] and Jacqueline H. Chen[a]

[a]Combustion Research Facility, Sandia National Laboratories, Livermore, CA, USA
[b]SINTEF Energy Research, Trondheim, Norway
[c]Ansaldo Energia, Baden, Switzerland

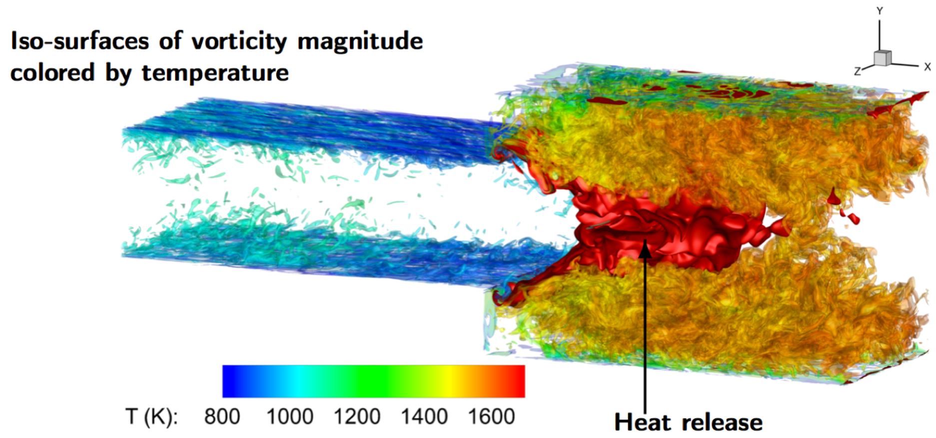

Acknowledgements:
Basic Energy Sciences, Office of Science, DOE
OLCF, NERSC, Norwegian CCS Research Centre (NCCS)

Aditya et al. PCI 2018

Staged gas turbine combustion

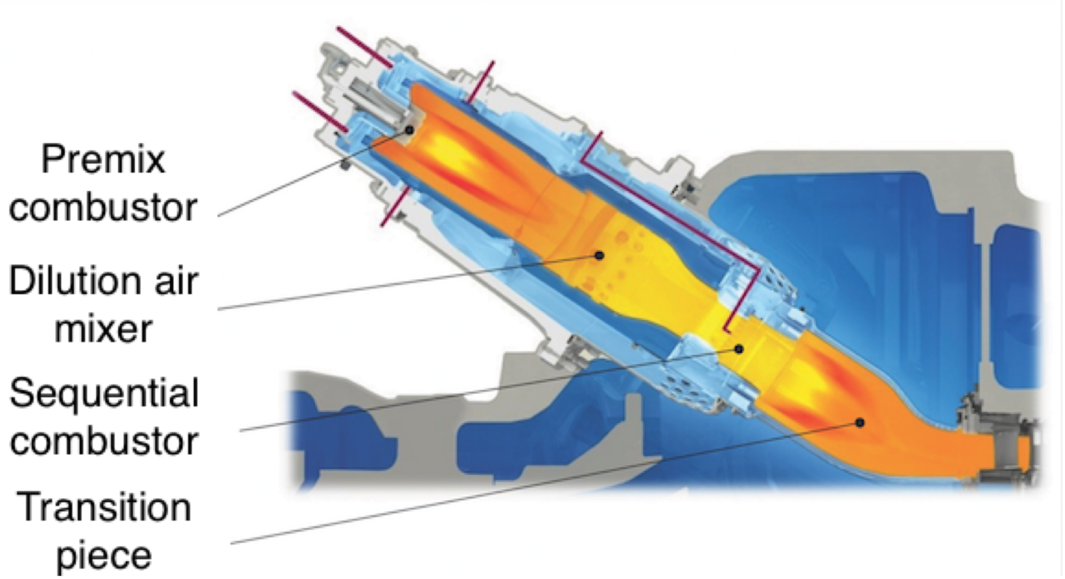

- Originally developed by ABB for high efficiency, load flexibility and low emissions
- Recently improved and simplified (reduced cost) for the H-class GT36 (538MW of power)
- First (premix) combustion stage based on flame propagation
- Second (sequential) combustion stage based on auto-ignition

Reheat burner

DNS of idealized reheat burner configuration from Ansaldo Energia

Operating conditions:

- Inlet temperature: ~ 1100 K
- Pressure: ~ 20 atm

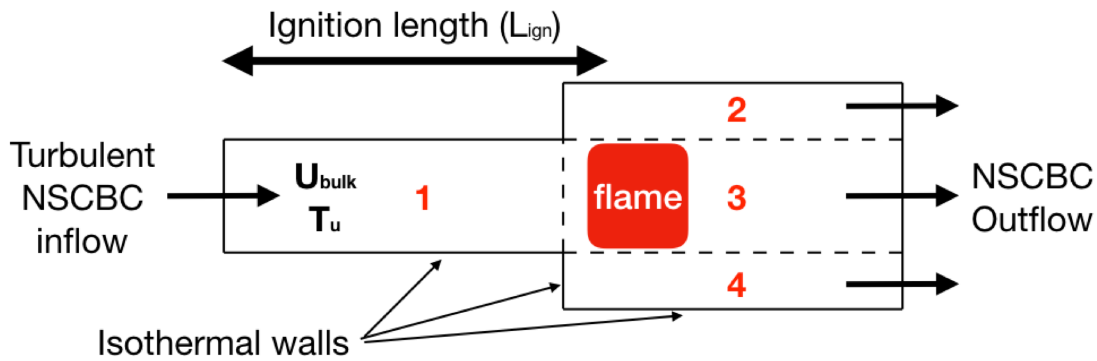

Scaled conditions:

- Mean inlet temperature:
- Pressure: 1 atm
- Fuel: hydrogen

Objective:

- Understand the flame stabilization
- Identify the modes of combustion
- Quantify the role of auto ignition

Simulation parameters

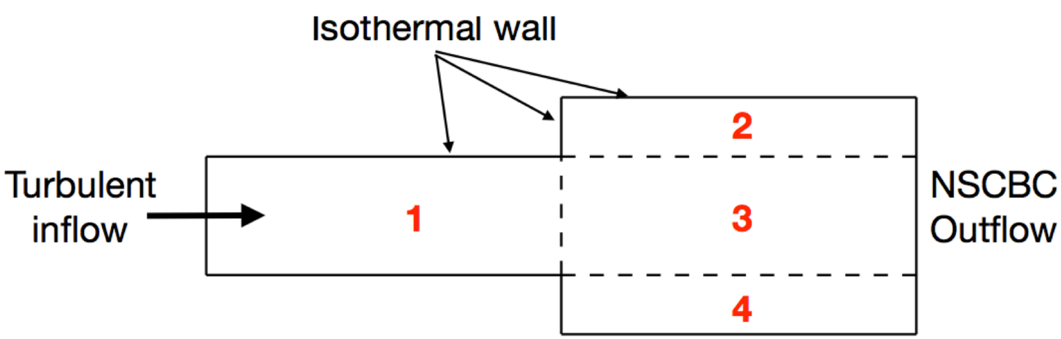

1.25 billion grid points
20 million CPU hours
$Re_b = 13000$

- **Chemical mechanism:** 9 species hydrogen-air (Li et al., 2004)
- **Inflow composition:** premixed H2 + O2 + N2 + H2O ($\phi = 0.35$)
- U_{bulk} = 200m/s, u′ = 20m/s, T_{inlet} = 1100K, T_{wall} = 750K
- **Inflow profile:** feed from DNS of a fully developed channel flow,

$Re_b \sim 13,000$, $Re_\tau = 378$ (viscous length scale)

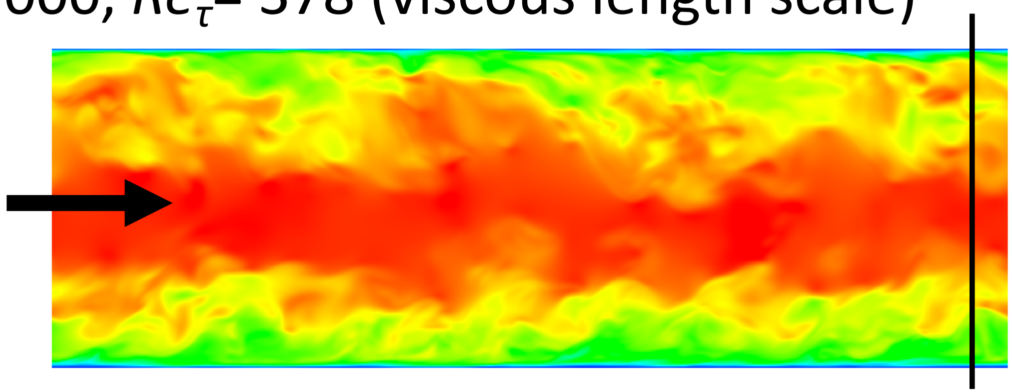

Feed data sampling plane

Two Combustion States

Iso-surfaces of vorticity magnitude colored by temperature

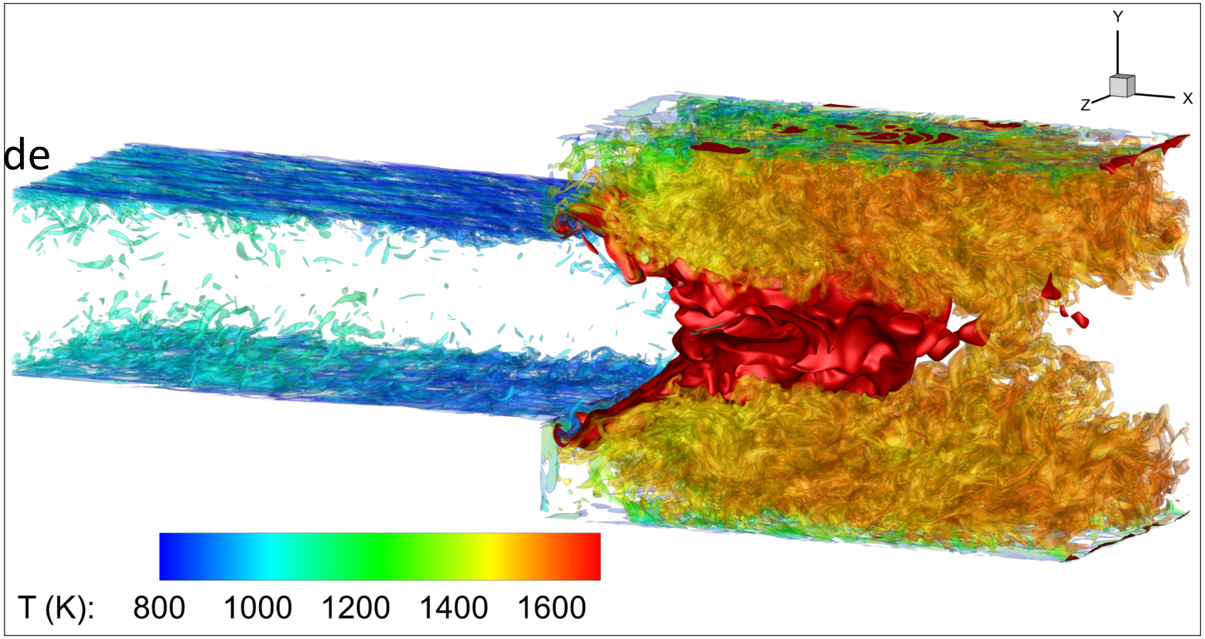

Two combustion configurations are observed:

- Design state: mainly auto-ignition in the combustion chamber
- Intermittent auto-ignition state: ignition in mixing section

Design combustion state

Heat release rate

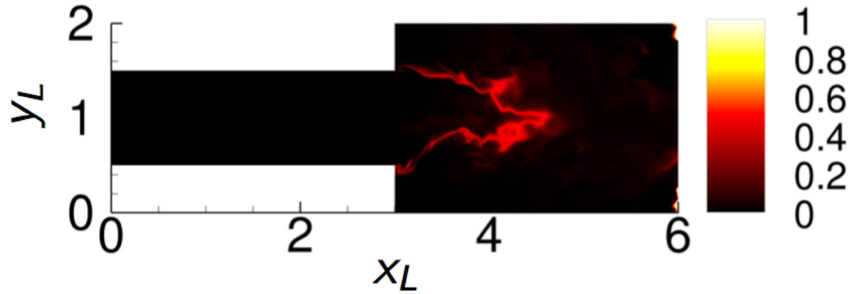

Temperature

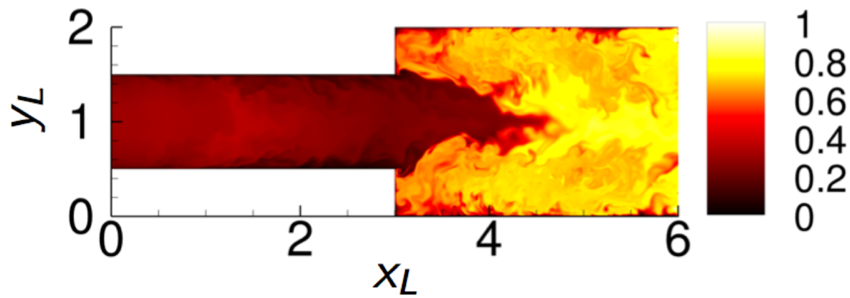

Mass fraction of HO_2

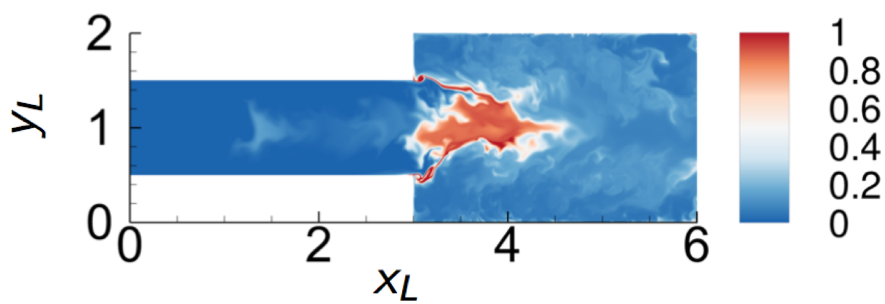

Combustion modes:

- Autoignition along center-line
- Flame propagation near corners
- HO2: indicative of chain branching

Transport budget analysis

$$\frac{\partial(\rho Y_{OH})}{\partial t} = -\nabla_\beta \cdot (\rho Y_{OH}\mathbf{u}_\beta) - \nabla_\beta \cdot (\rho Y_{OH}\mathbf{V}_{\beta,OH}) + W_{OH}\dot{\omega}_{OH}$$

Advection Diffusion Reaction

- Auto-ignition: balance between advection and reaction
- Flame propagation: balance between diffusion and reaction

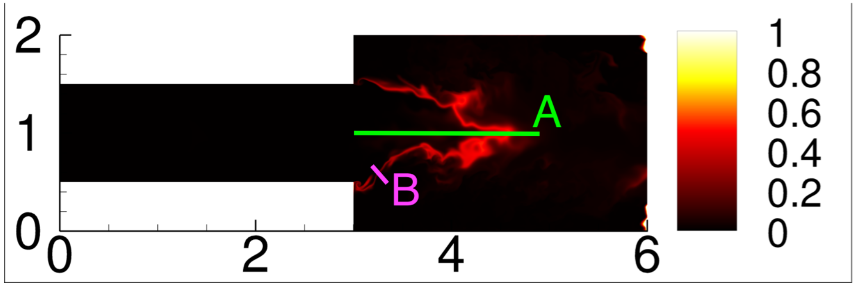

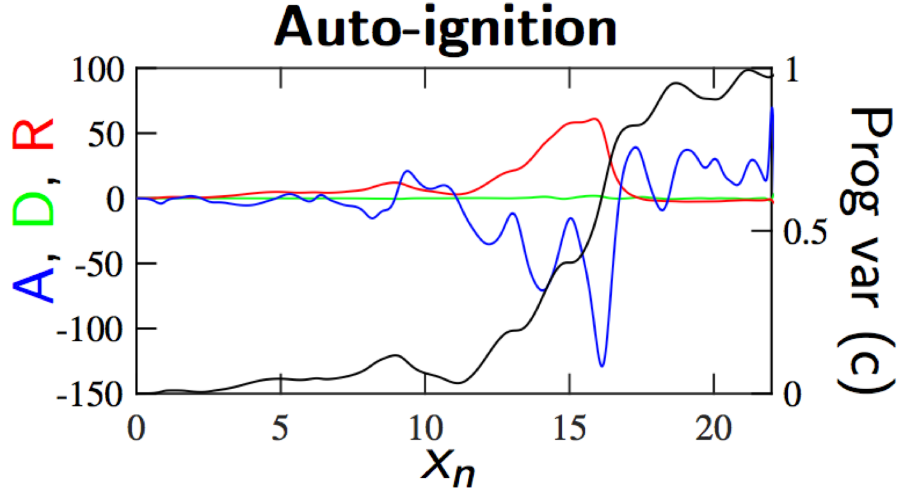

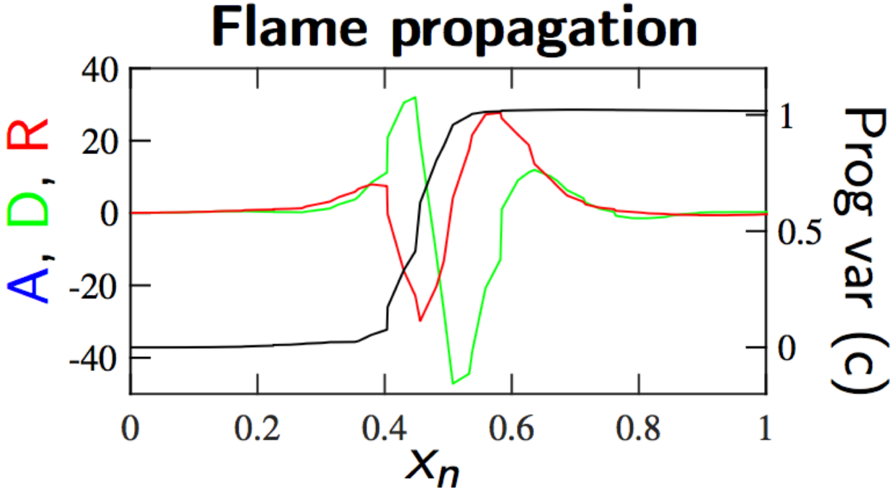

Modes segmented with chemical explosive mode analysis (CEMA)

- $\alpha = \varphi_s / \varphi_\omega$: ratio of the projected non-chemical source term and the projected chemical source term (C. Xu et al., PROCI 2018)

Three modes are identified:

- Assisted-ignition ($\alpha > 1$): diffusion significantly promotes reaction
- Auto-ignition ($-1 < \alpha < 1$): chemistry plays a dominant role
- Extinction zone ($\alpha < -1$): diffusion dominates chemistry and suppresses ignition

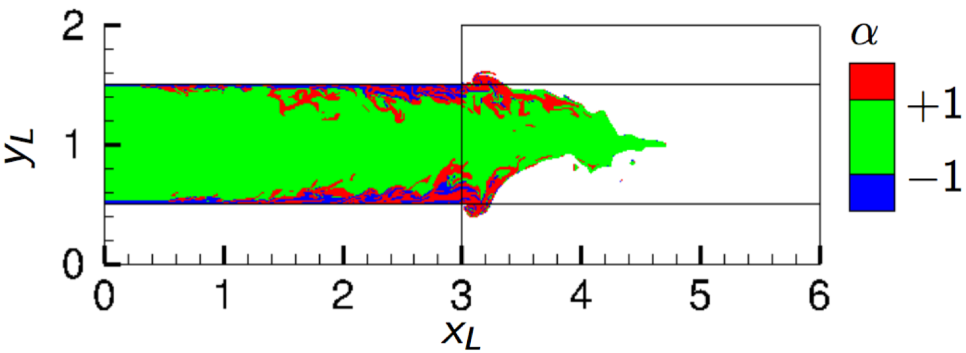

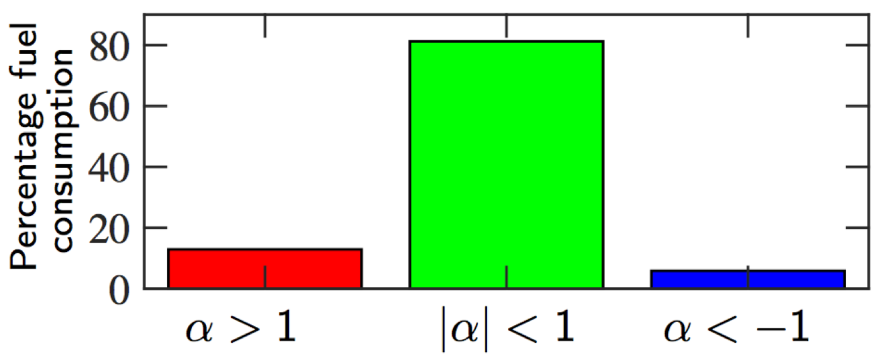

Intermittent auto-ignition state

Heat release rate

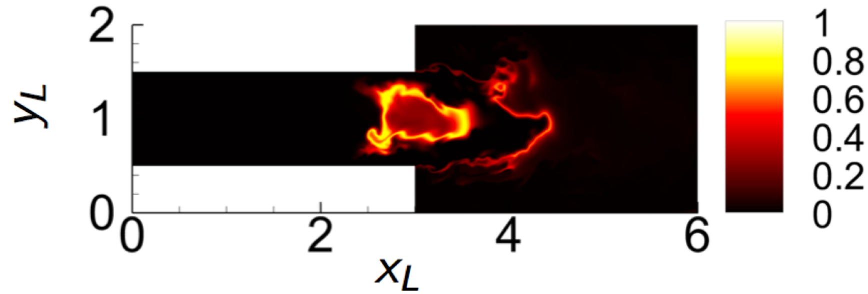

Temperature

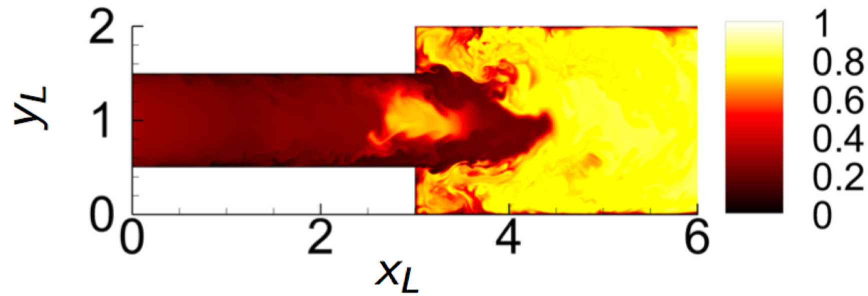

Mass fraction of HO_2

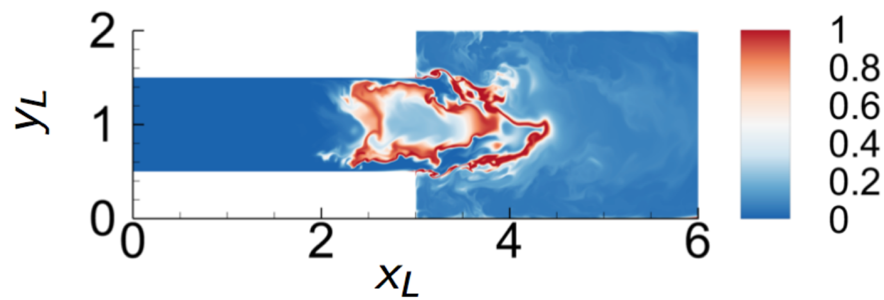

- Early auto-ignition in the mixing section
- Ignition kernel advects downstream
- Occurs intermittently

Intermittent auto-ignition state

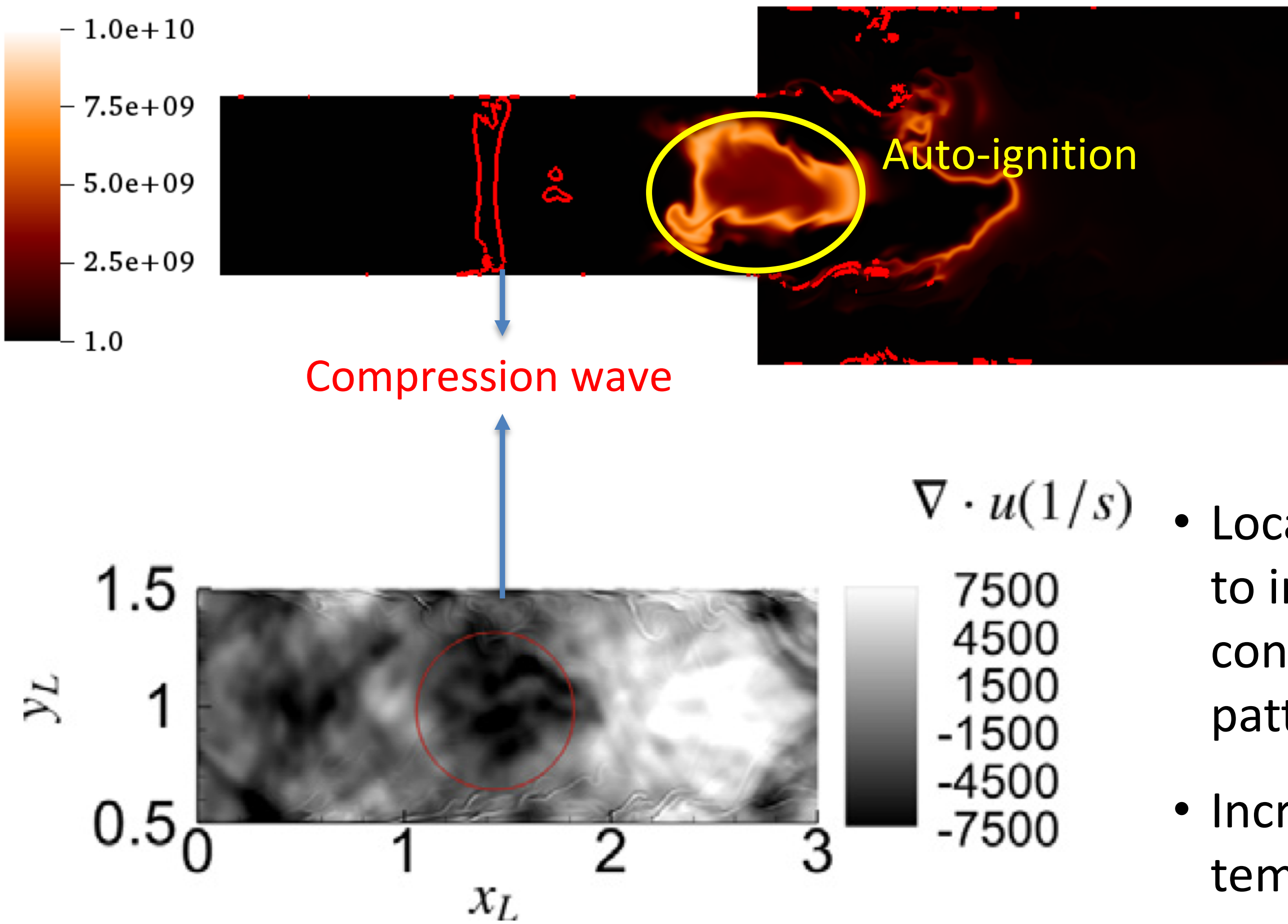

Contours of heat release

Contours of instantaneous dilatation

- Local rise in pressure due to intermittent constructive interference pattern
- Increases local temperature by 20-30 K
- High reactivity of hydrogen
- Decrease in ignition delay time (30%)

Conclusions

- Performed DNS of a reheat burner at scaled conditions
- Two states of hydrogen/air combustion have been observed:
  - design state: flame propagation and auto-ignition in the combustor
  - intermittent auto-ignition in mixing section
- Premature auto-ignition arises due to pressure (and following temperature) rise in mixing section
- Quantified the contribution of different modes towards heat release using chemically explosive mode analysis (CEMA)

Constraints imposed by exascale architecture

- **Power:** primary design constraint for future HPC system design
- **Cost:** Data movement dominates: optimize to minimize data movement
- **Concurrency:** Exponential growth of parallelism within chips
- **Locality***:* must reason about data locality and possibly topology
- **Memory Scaling**: Compute growing 2x faster than capacity or bandwidth, **Heterogeneity:**Architectural and performance nonuniformity

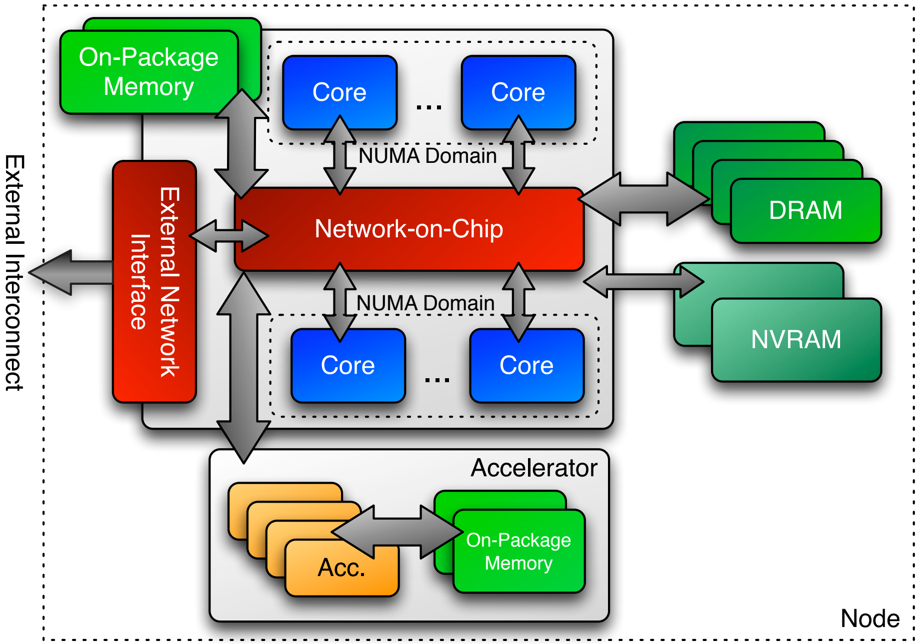

Conceptual model of future HPC node

Express data locality and independence, express massive parallelism, minimize data movement and reduce synchronization, detect and address faults

Parallel Programming 101 - Productivity

Functionally correct application code

Mapping to target machine

Extraction of para

anagement of
ata transfers

Task scheduling
Latency hiding

Data-Dependent
Behavior

Parallel Programming 101

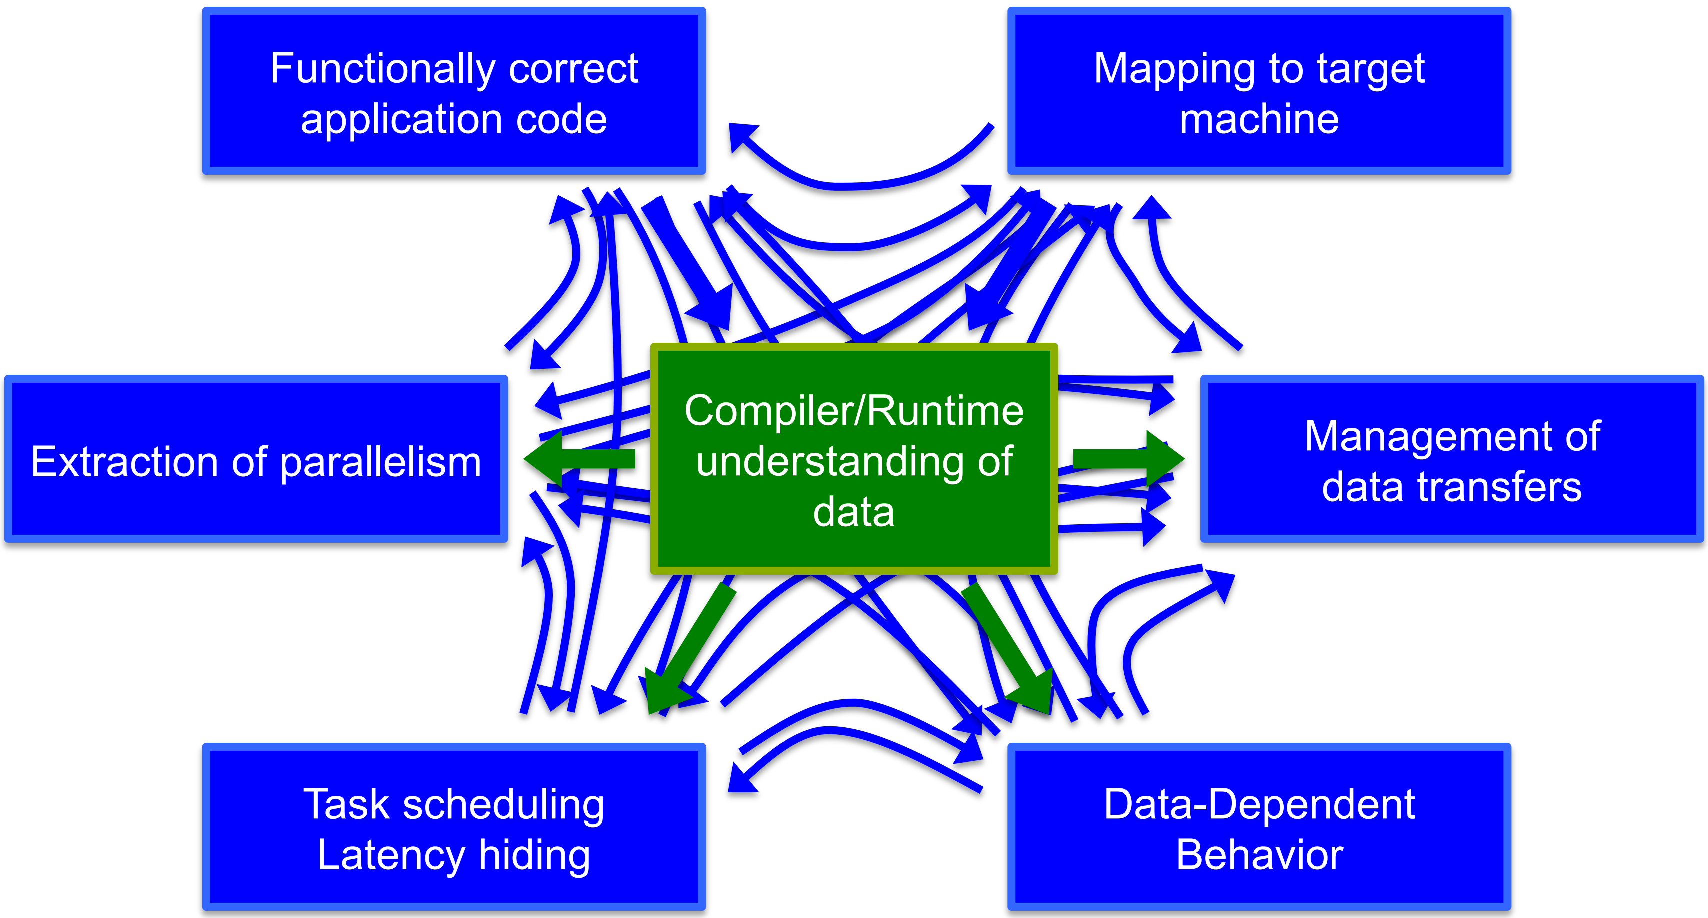

Legion programming system applied to S3D

- A data-centric parallel programming system
- A programming model for heterogeneous, distributed machines
  - Automates many aspects of achieving high performance, such as extracting task- and data-level parallelism
  - Automates details of scheduling tasks and data movement (*performance optimization*)
  - Separates the specification of tasks and data from the mapping onto a machine (*performance portability*)
- Legion application example: S3D
  - Production combustion simulation
  - Written in ~200K lines of Fortran
  - Direct numerical simulation using explicit methods

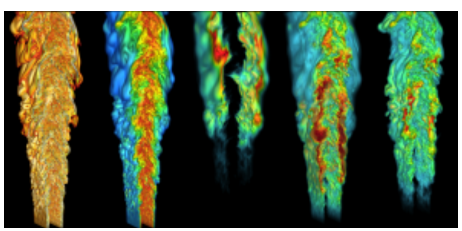

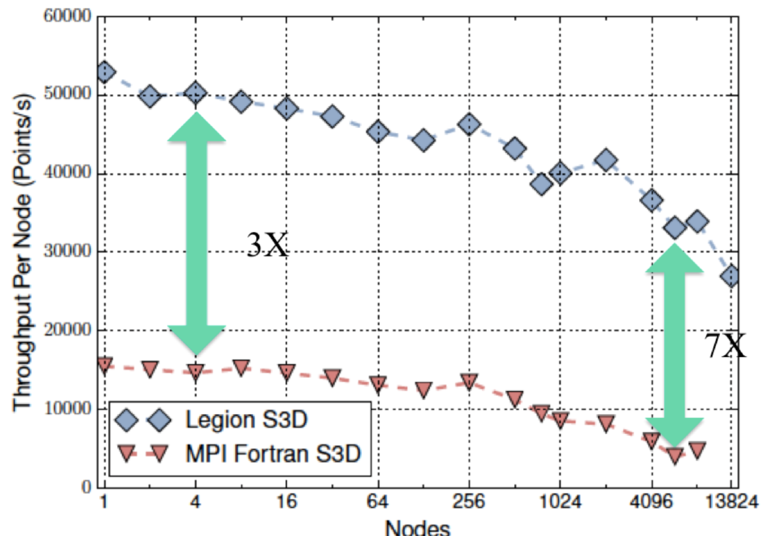

S3D performance Legion vs. MPI

S. Treichler et al., "S3D-Legion: An Exascale Software for Direct Numerical Simulation (DNS) of Turbulent Combustion with Complex Multicomponent Chemistry," CRC Book on Exascale Scientific Applications: Programming Approaches for Scalability Performance and Portability, 2017.

Legion task graph for 1 time step for S3D-Legion DNS running ndodecane (35 species) on a single Summit node

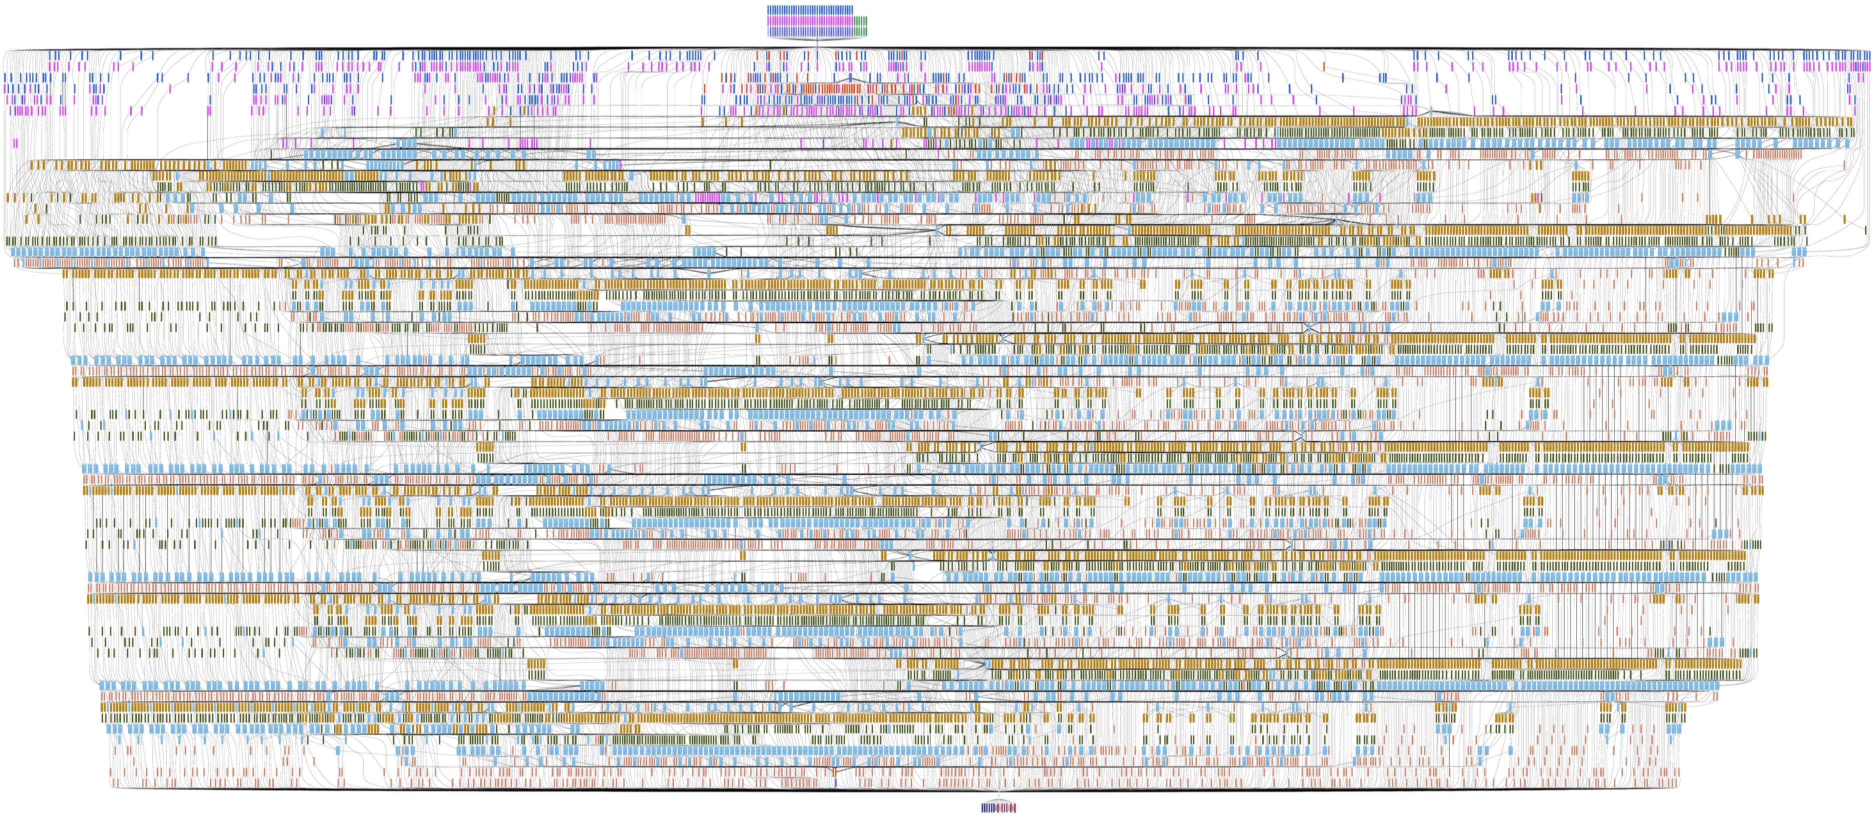

- Nodes – tasks (compute kernels or data copies)
- Edges – ordering dependencies between the nodes

Exploring the design space of in situ workflows: constraints and observations

- Minimize performance impact to the simulation
  - Work within time and memory constraints
  - Minimize cache impact

- Behavior of analysis algorithms varies widely
  - Data dependencies, communication patterns, scalability, instruction mixes, time and memory requirements
  - Data dependent algorithms are very hard to characterize

Analytics Algorithms Used by Combustion Applications Differ in Behaviors (from Solvers and with Each Other)

Build graphs
Integer ops
Branching
Data dependent

Floating point
Compute intensive
Data-parallel
Low communication

- Multi-variate volume and particle rendering
- Particle querying and analysis
- Topological feature segmentation:
  - Contour trees
  - Morse-Smale complex
  - Time tracking
  - Scalar field comparison
- Filtering and averaging
- Shape analysis
- Statistical moments (conditional), dimensionality reduction (joint PDFS), spectra (scalar, velocity, …)
- UQ driven analytics (quantities of interest)
- Machine Learning and anomaly detection

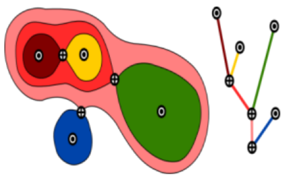

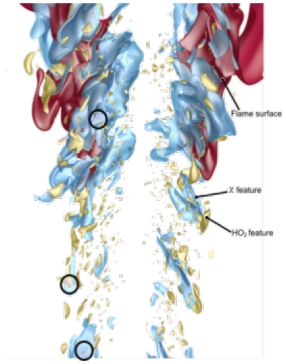

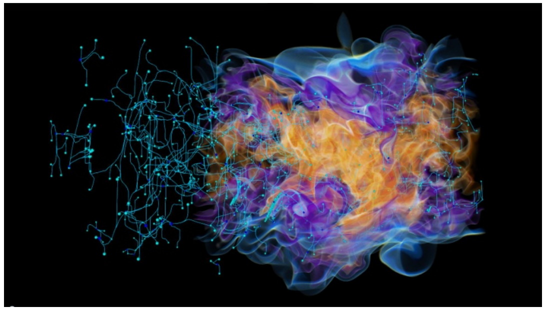

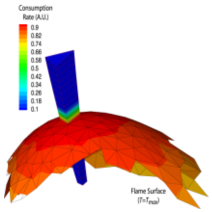

Flame-centric control volume analysis

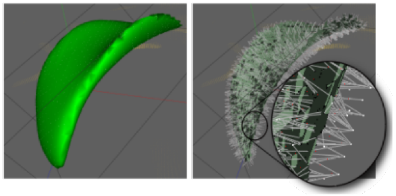

In-situ data analytics in Legion
Chemical Explosive Mode Analytics (CEMA)

- CEMA: eigenvalue solve on the reaction rate Jacobian to determine the mode of combustion

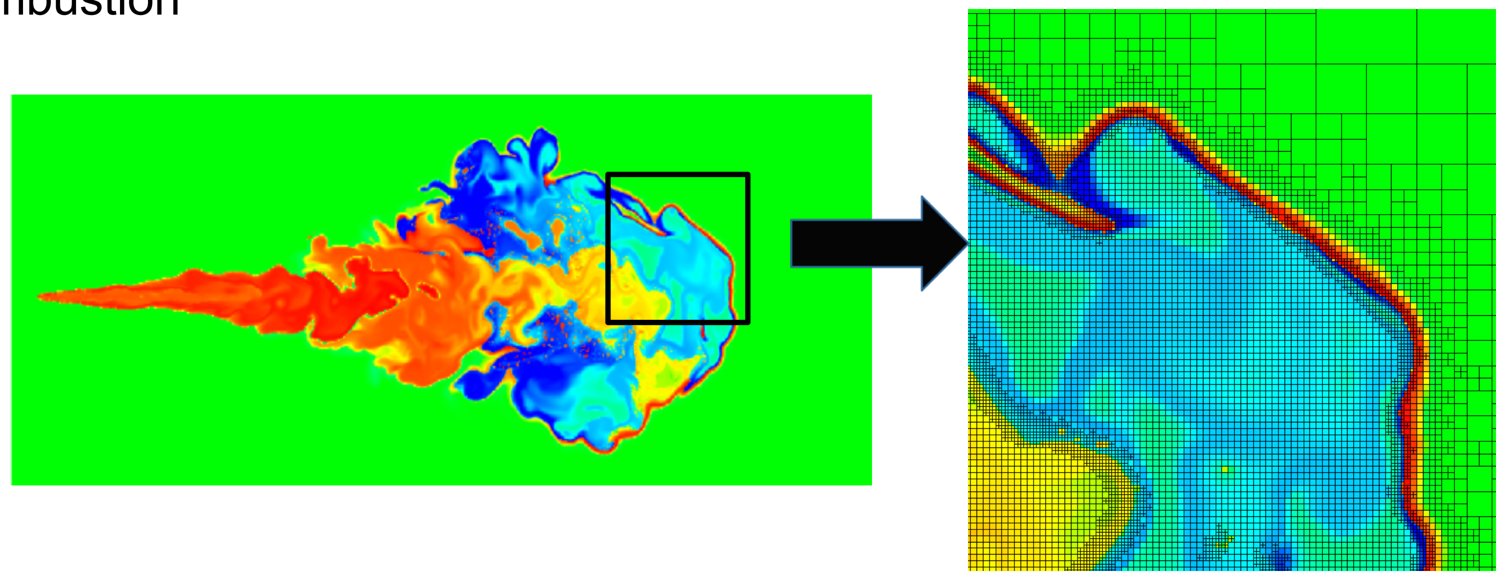

- Run CEMA at each time step as a diagnostic to steer mesh refinement
- CEMA computation takes longer than a single explicit RK stage (6 stages/timestep)
- Dividing CEMA across RK stages and interleaving with other computation so as not to impact other critical operations would be hard to schedule manually
- Asynchronous task execution, schedule CEMA on CPU resources
- Interoperate Fortran CEMA with Legion code – took a day to implement

Execution overhead of in-situ analytics (CEMA) in S3D-Legion is small (Titan & Piz Daint)

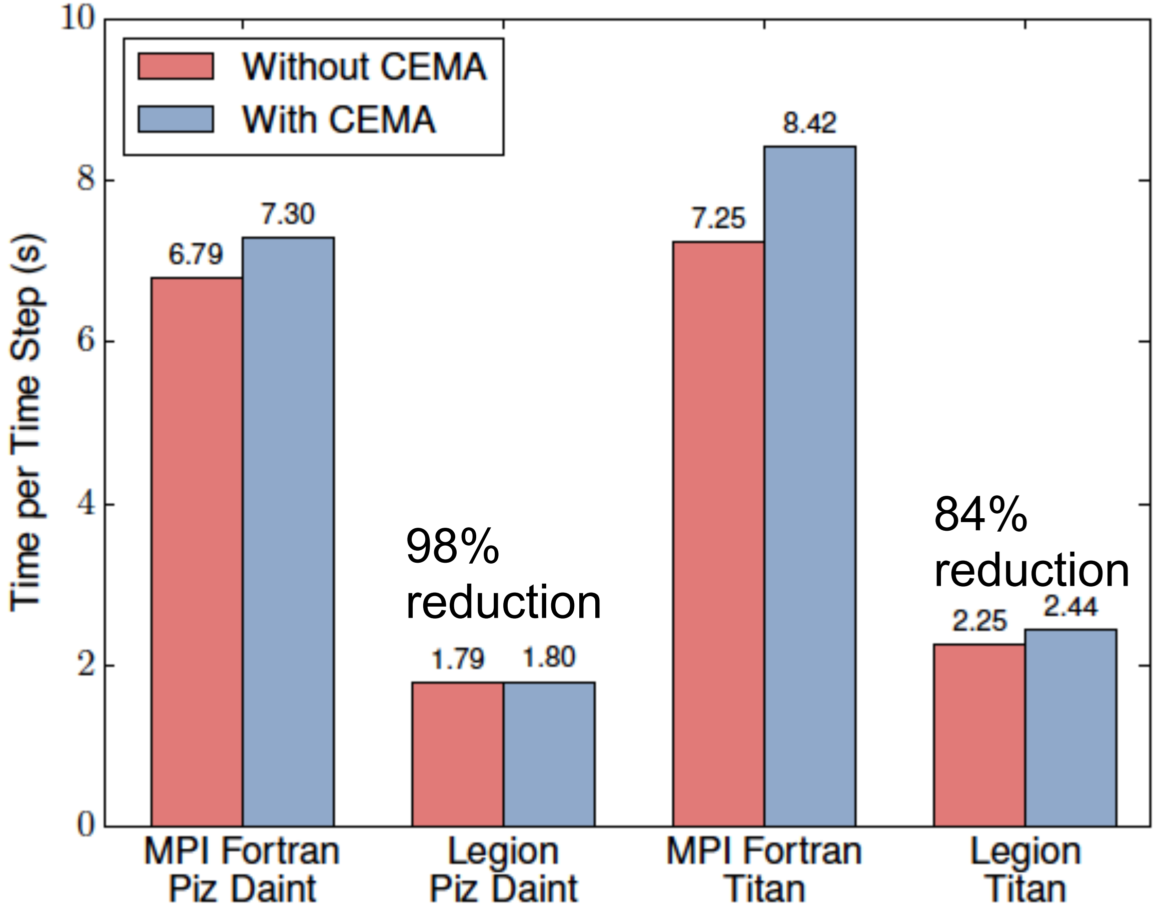

Legion S3D lessons learned

- Legion
  - S3D shows potential of data-centric, task-based models
  - Enables new simulation capabilities (physics, and in situ analytics)
  - Code is easier to modify and maintain
    - Ports are just new mappings, easy to tune for performance
    - New functionality usually just means new tasks
    - Legion will figure out the dependences and scheduling
    - Strong scaling well enables more complex workflows with analytics
    - Productivity requires higher level abstraction layer for scientists to write in

- Co-Design and ECP
  - The Legion/S3D experience is a tribute to application co-design
  - Computer and computational scientists worked closely
  - Major progress on important problems resulted

Exascale Targets: Science at Relevant Engine Conditions

- **Hybrid DNS/LES (near DNS) with dynamic adaptive mesh refinement, multi-physics (sprays, soot, radiation at high pressure) in geometry**
- **Reactivity Stratified Compression Ignition IC Engines** - multi-stage, high pressure autoignition of a liquid hydrocarbon fuel blend, soot formation
- **Natural Gas IC Engines – ignition and knock**
- **Scramjets** – cavity stabilized shear driven lean turbulent premixed flames, effect of products recirculation coupled with high Re, high Ka, compressible flames
- **Gas Turbines** – swirl stabilized spray combustion gas turbines with lean premixed combustion, flame stabilization, nitric oxide emissions, thermo-acoustics
- **Include in-situ analytics & visualization**