SAND2019-7676C

DNS of Turbulent Combustion

Jacqueline H. Chen

Sandia National Laboratories

jhchen@sandia.gov

Cambridge Combustion

Institute Summer Shool

July 8-12, 2019

Office of Basic Energy Sciences — Serving the Present, Shaping the Future

Office of Science
U.S. Department of Energy

Office of Science U.S. Department of Energy
INCITE
Advancing America's Science and Industrial Competitiveness

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Motivation, DNS Governing Equations, Numerical Methods and HPC Considerations

Jacqueline H. Chen

Sandia National Laboratories

jhchen@sandia.gov

Cambridge Combustion Institute Summer Shool

July 8-12, 2019

Combustion and Technology Advancement

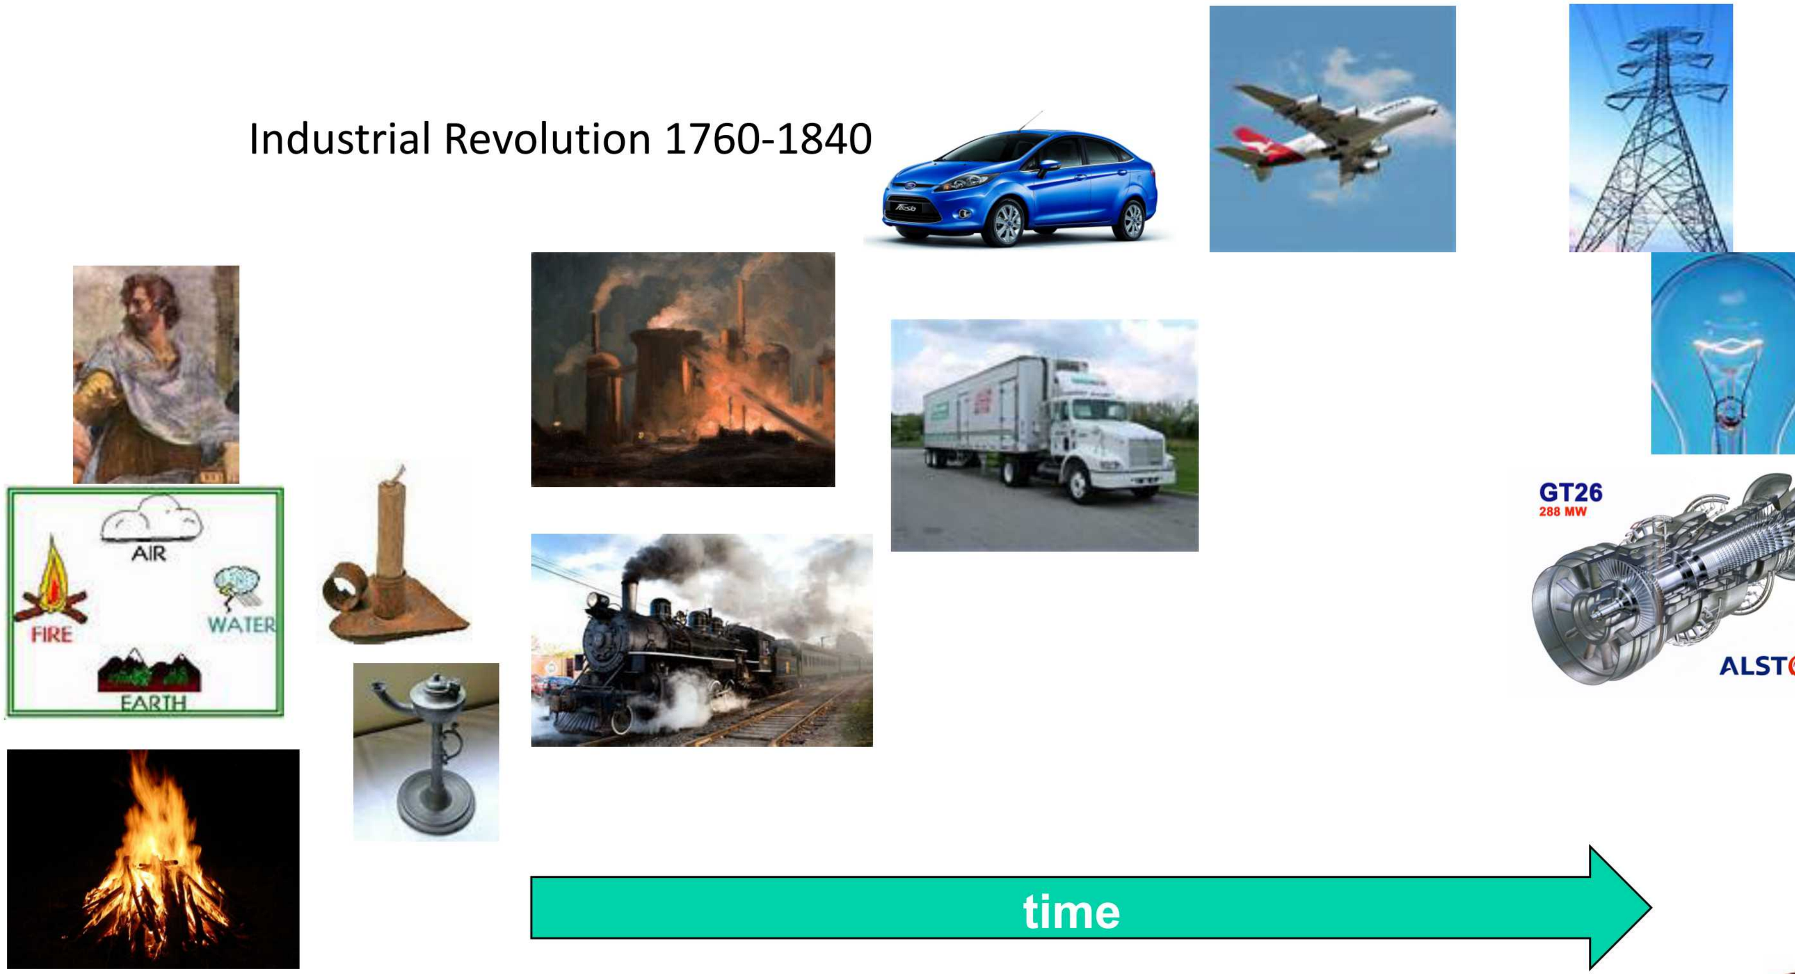

Energy consumption – where does it all go?

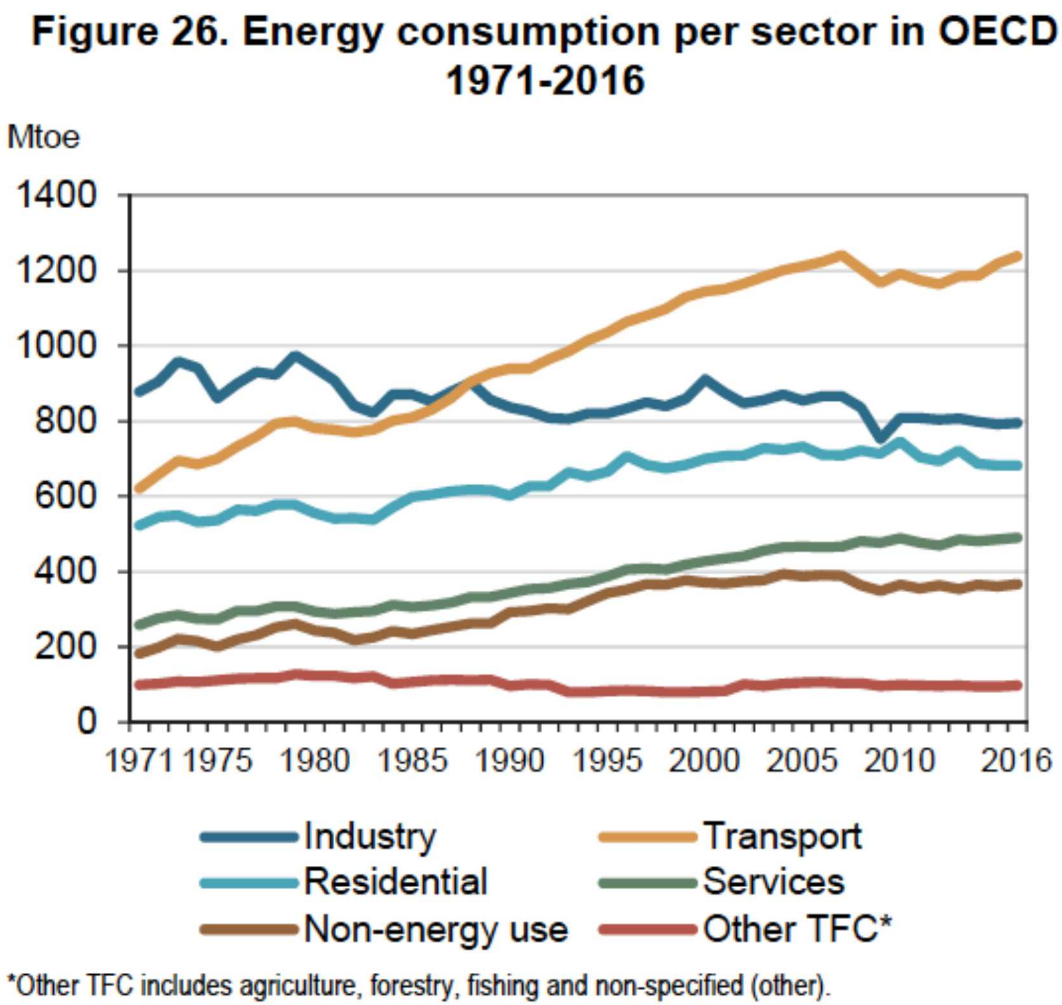

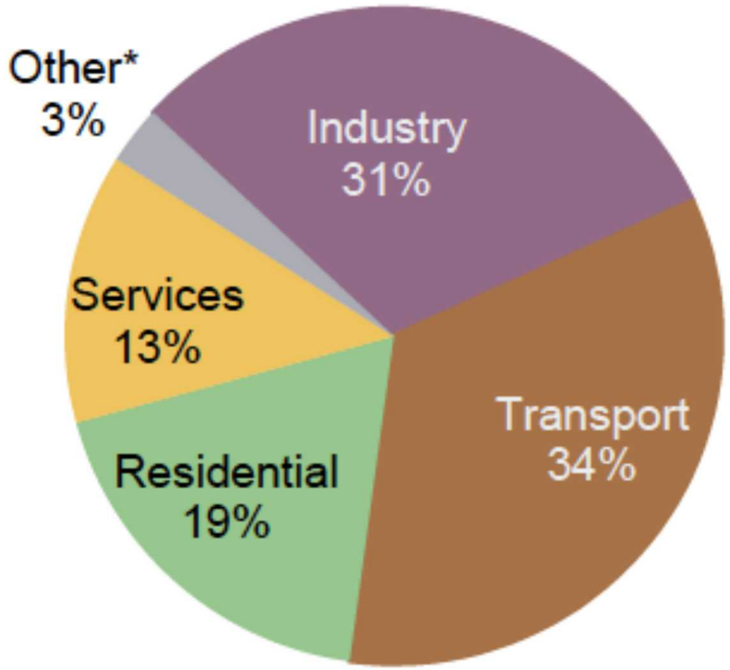

IEA World Energy Balances 2018

Source of Transportation World Energy Consumption

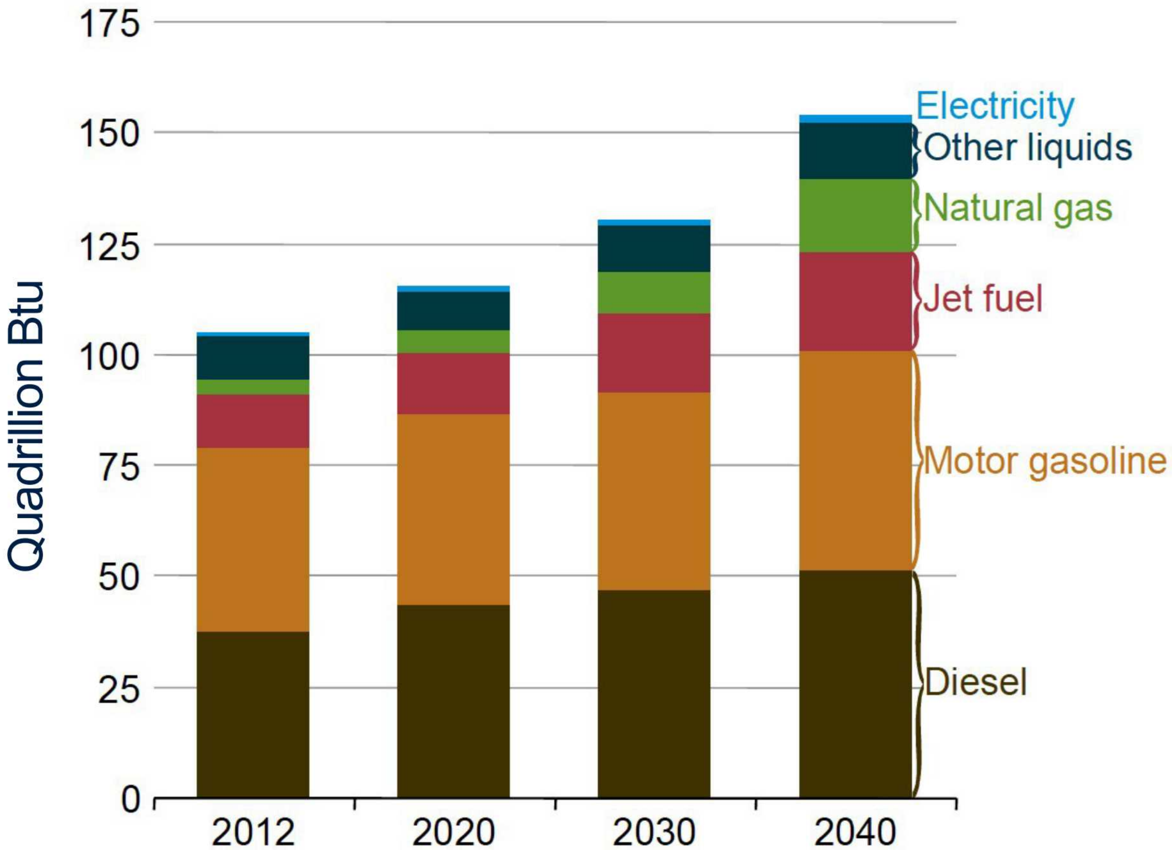

DOE IEA Oulook 2016

Source of World Net Electricity Generation

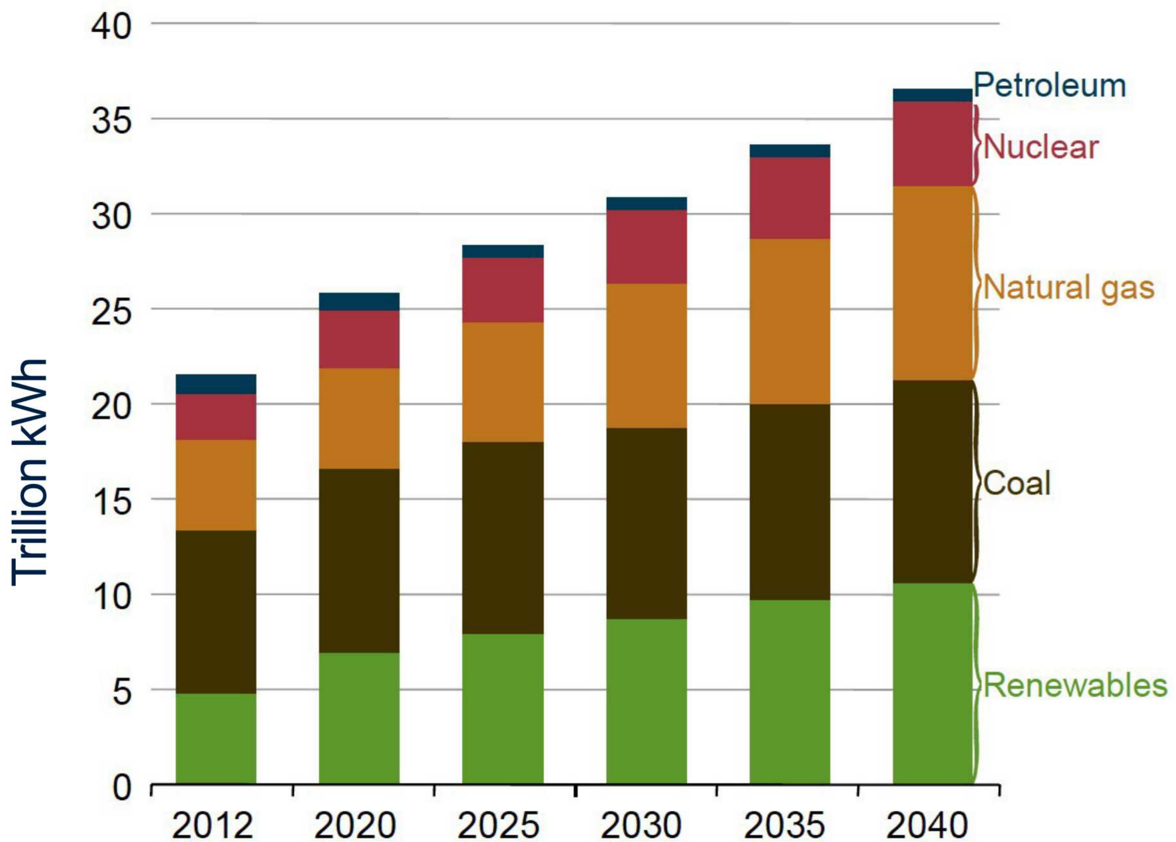

DOE IEA Oulook 2016

Temperature's rising

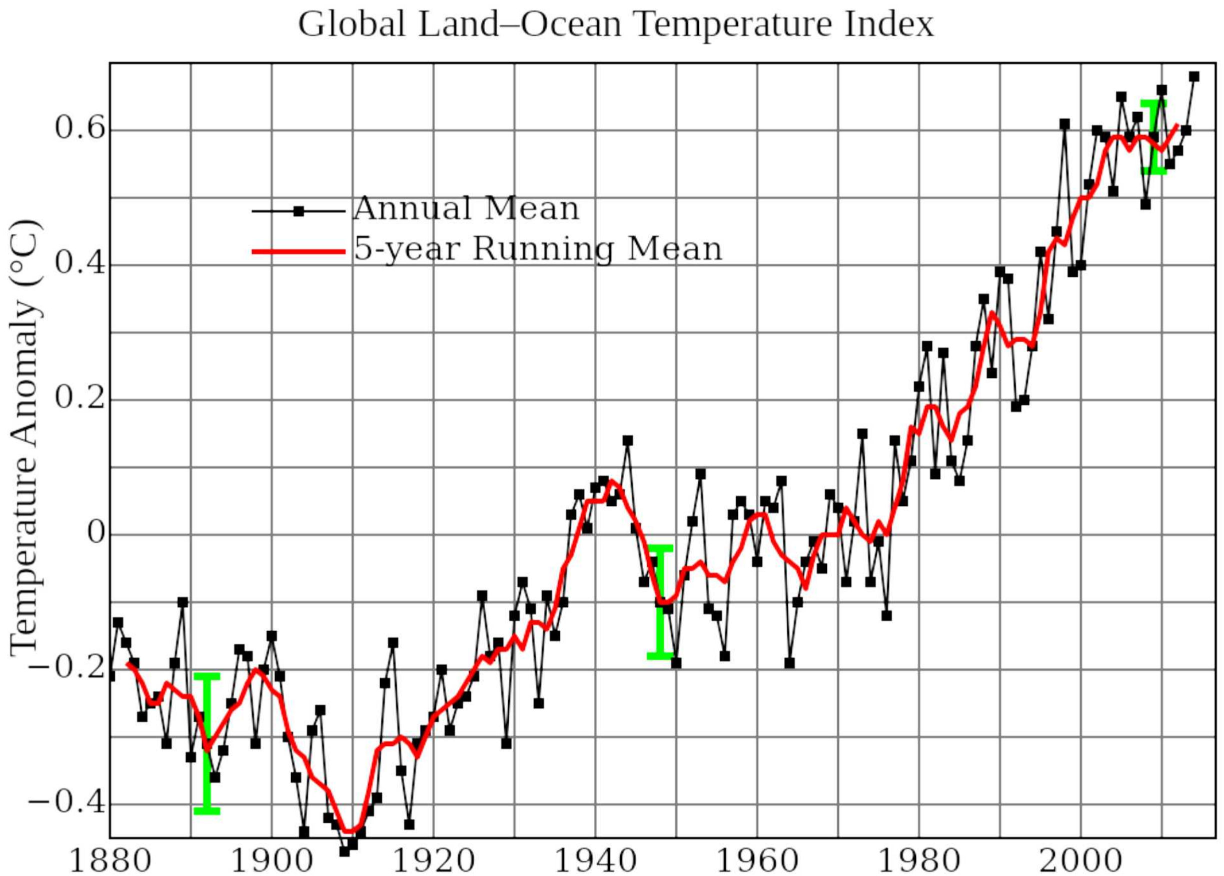

Global mean surface temperature change from 1880 to 2014, relative to the 1951–1980 mean. The black line is the annual mean and the red line is the 5-year running mean. The green bars show uncertainty estimates. Source: NASA GISS.

Where are the greenhouses coming from?

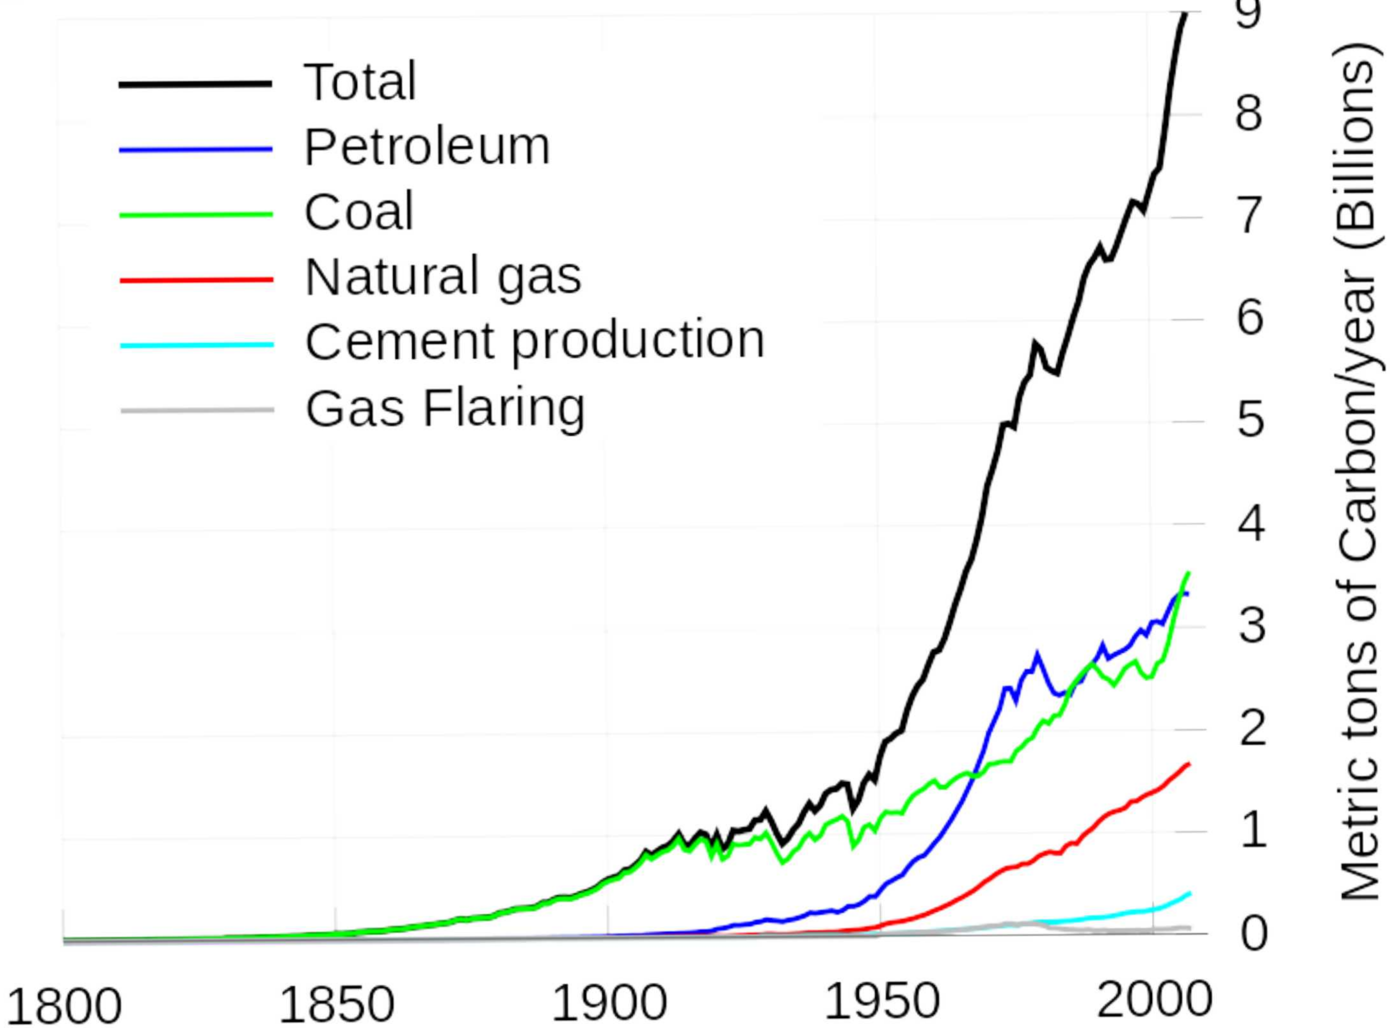

Hard to beat the energy density of petroleum!

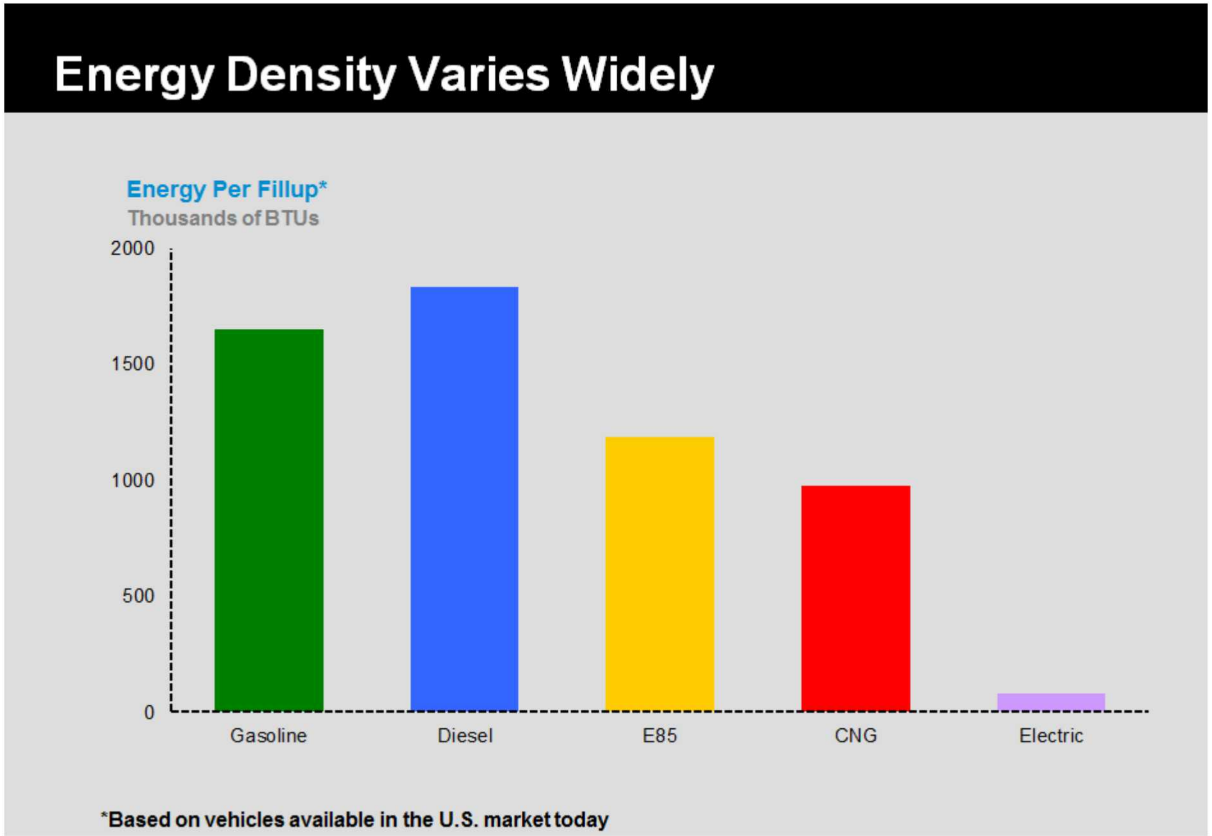

All of the energy concentrated in one gallon of gasoline is enough to charge an iPhone once a day for almost 20 years.

Mitigation Strategies for Greenhouse Gas Emissions

- CO_2-sequestration (Carbon Capture and Storage, CCS)
- Bio-fuels. (Ethanol)
- Hydrogen-rich fuels
- Increase fuel efficiency
- Other….

→ Sound combustion foundation for technology development

Combustion is Multi-physics & Multi-scale

CN45, Glow Plug Off

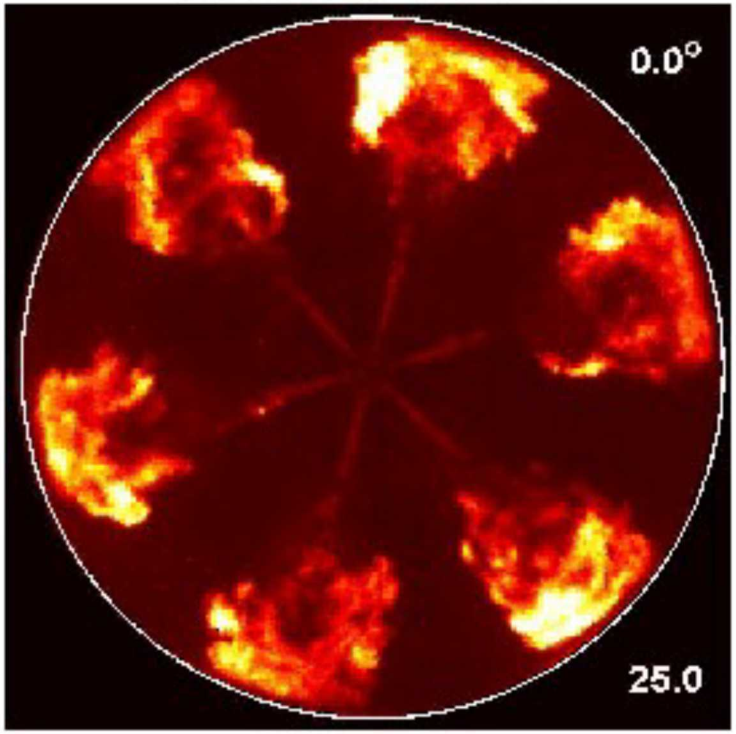

Diesel Engine Autoignition, Soot Incandescence
Chuck Mueller, Sandia National Laboratories

Large range of length and time scales

- In-cylinder geometry (cm)
- Turbulence-chemistry (microns-mm)

–Soot inception (nanometer)

Chemical complexity

- large number of species and reactions (100's of species, thousands of reactions)

Multi-Physics complexity

- multiphase (sprays, gas phase, soot)
- thermal radiation

All these are tightly coupled

Motivation: Changing World of Fuels and Engines

- Fuel streams are rapidly evolving
  - Heavy hydrocarbons
    - Oil sands
    - Oil shale
    - Coal
  - New renewable fuel sources
    - Ethanol
    - Bio-diesel
    - Bio-gas
- New engine technologies
    - Direct Injection (DI)
    - Homogeneous Charge Compression Ignition (HCCI)
    - Low-temperature combustion
- Mixed modes of combustion
- Predictive, validated multi-scale models!

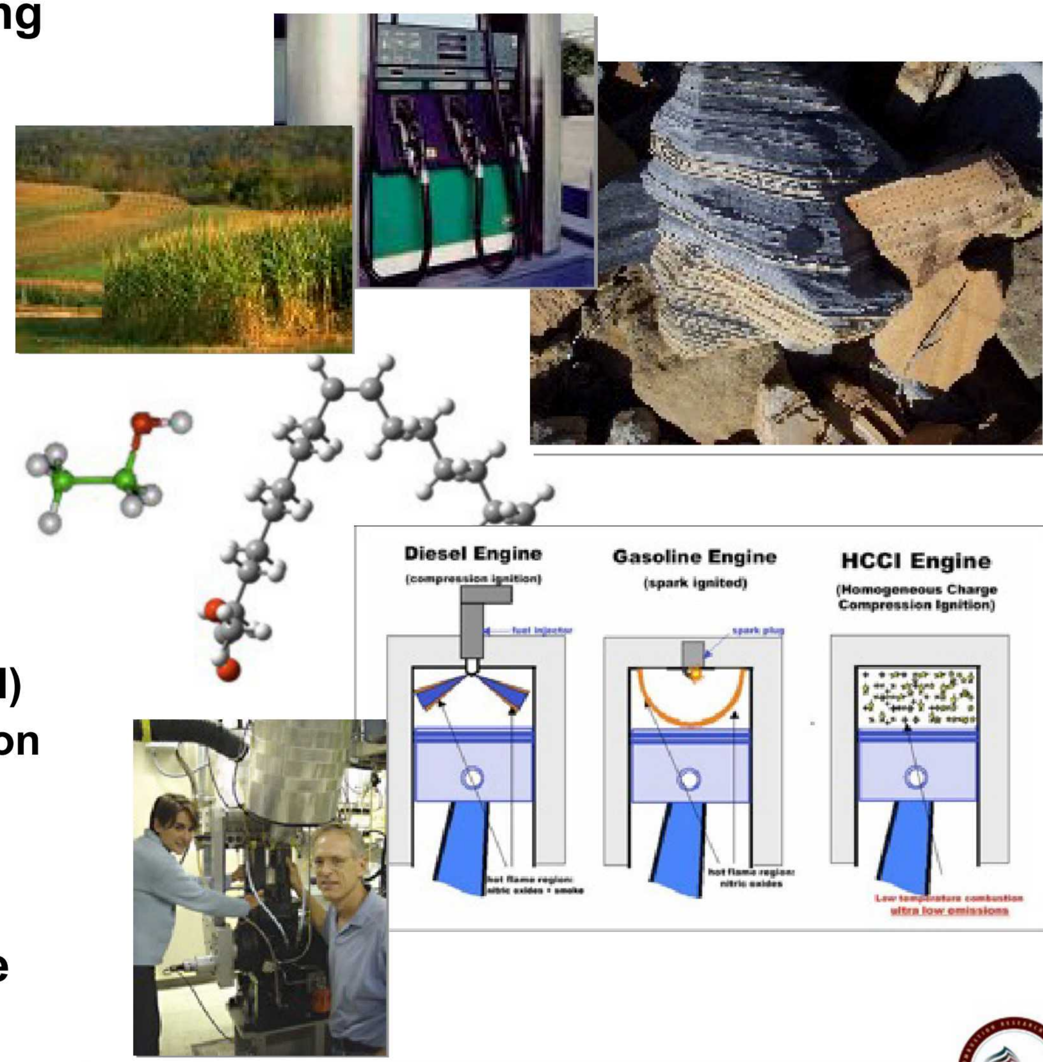

Comparison of Engines

Gasoline Engine
(Spark Ignition)

Diesel Engine
(Compression Ignition)

HCCI Engine
(Homogeneous Charge Compression Ignition)

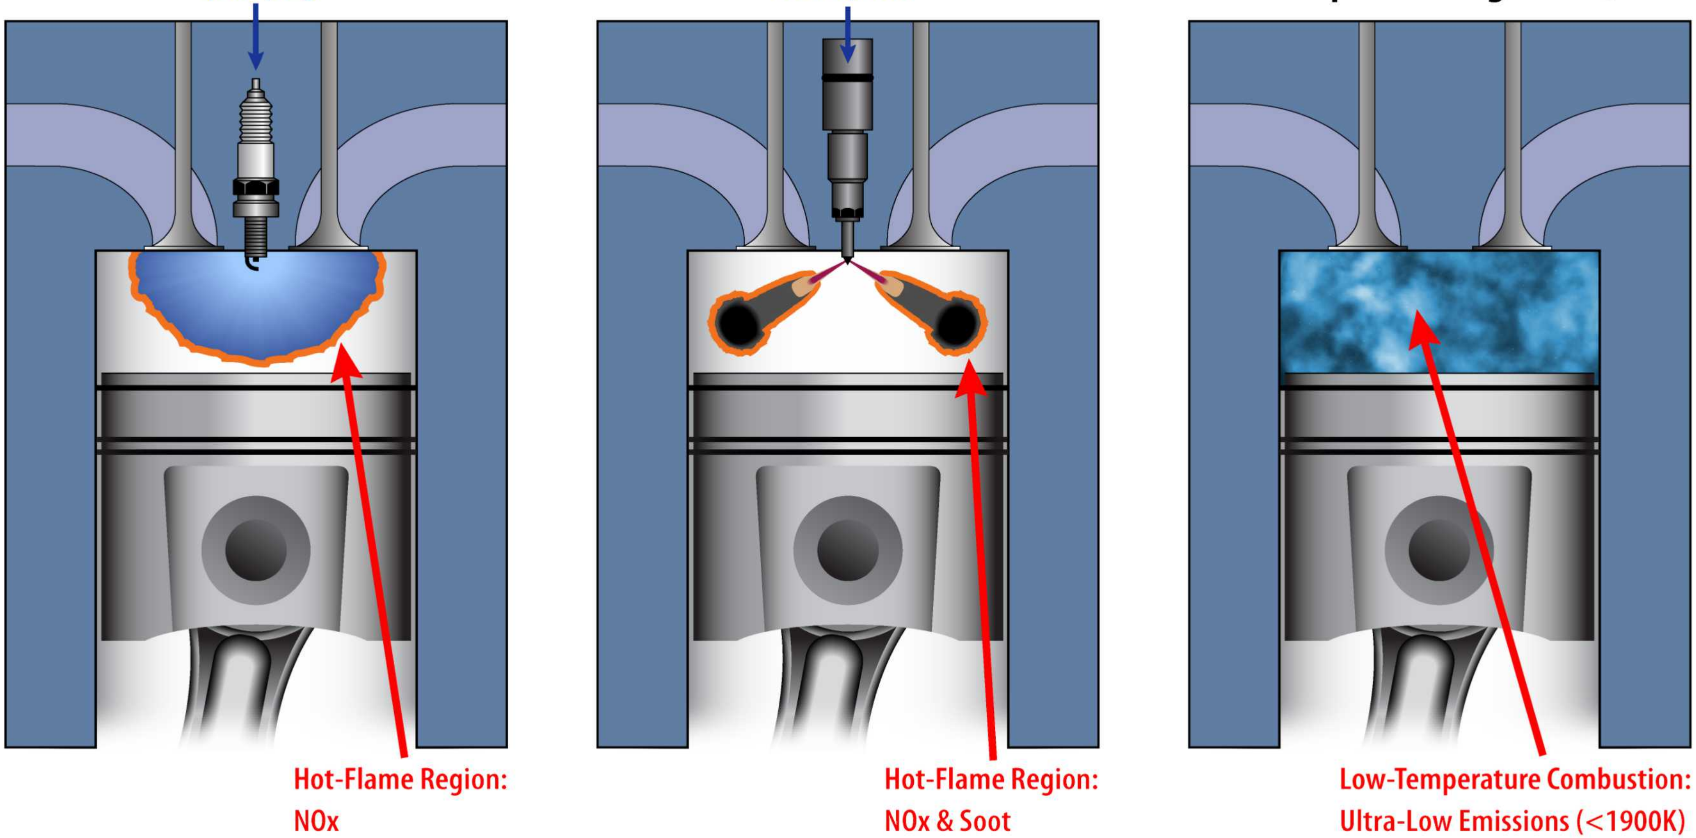

Fuel- and Load-Flexible Power Generation Towards Carbon Free Renewable Electricity

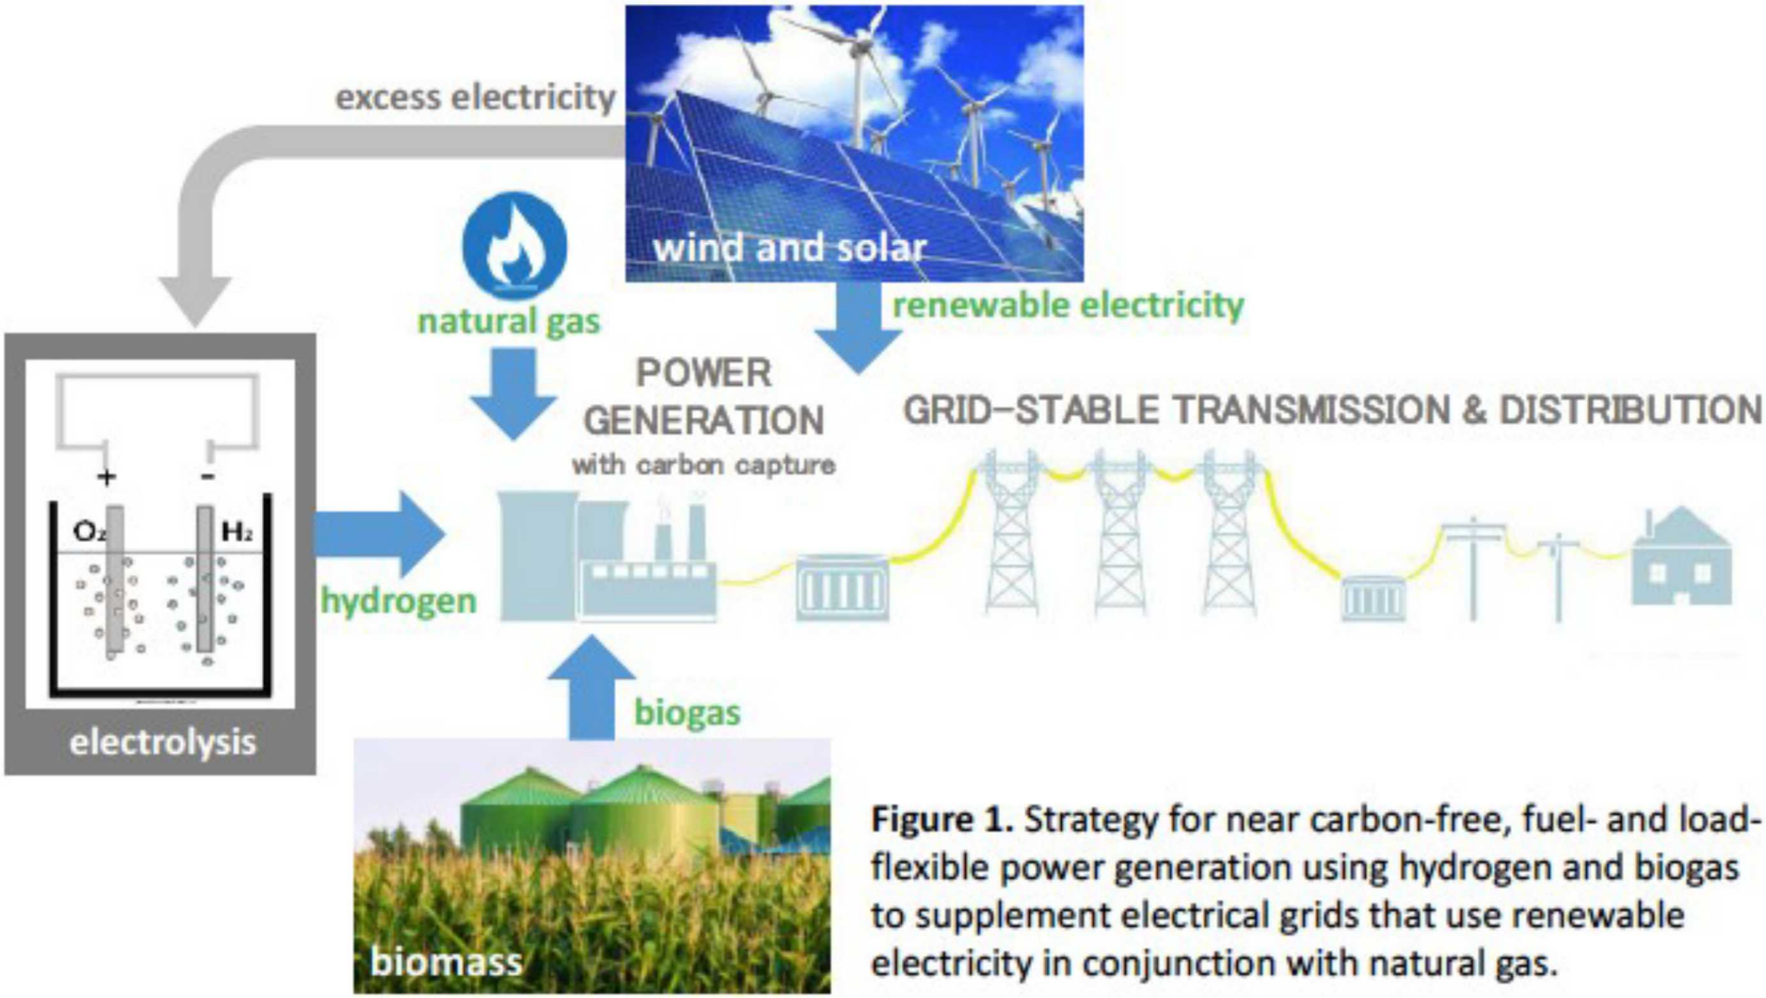

Figure 1. Strategy for near carbon-free, fuel- and load-flexible power generation using hydrogen and biogas to supplement electrical grids that use renewable electricity in conjunction with natural gas.

Courtesy Y. Ju

Challenges With Hydrogen-Rich Combustion

- Optimized hydrogen fuel transport and combustion system will enable safe operation within a broad range of fuels and operating conditions - enable fast fuel switches or blending (e.g. between natural gas and hydrogen-rich syngas) during fast start-ups and load-shifts.

- Compared to natural gas, hydrogen has higher stoichiometric temperatures, reduced ignition delay, smaller calorific value, requires higher volumetric flow rates to avoid preignition and to ensure flashback into premixer section doesn't occur

- Fundamental challenges to lean premixed hydrogen-enriched combustion in gas turbines:
  - F**lame flashback**
  - **Fuel injection: flame holding and stabilization**
  - **Thermo-acoustics**

Fuel and Load Flexibility Through Staging

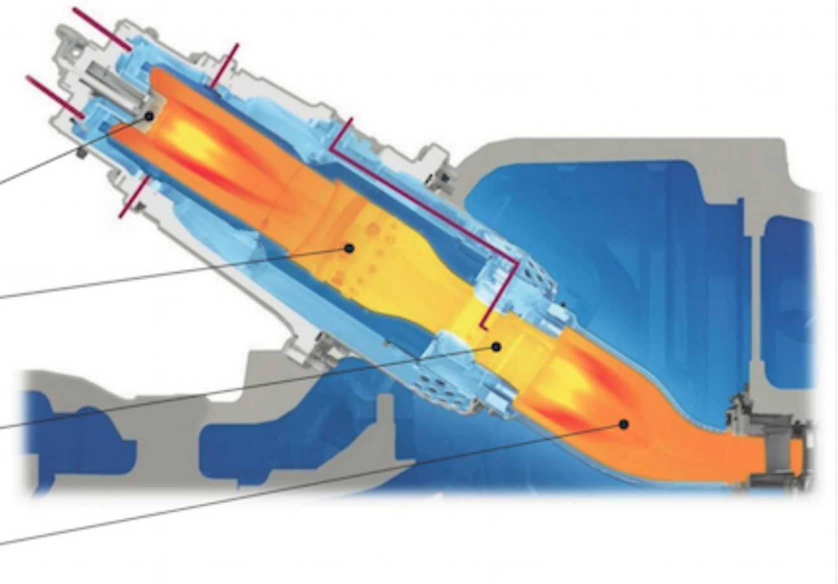

- Originally developed by ABB for high efficiency, load flexibility and low emissions
- Recently improved and simplified (reduced cost) for the H-class GT36
- First (premix) combustion stage based on flame propagation
- Second (sequential) combustion stage based on auto-ignition

Fuel and Load Flexibility Through Staging

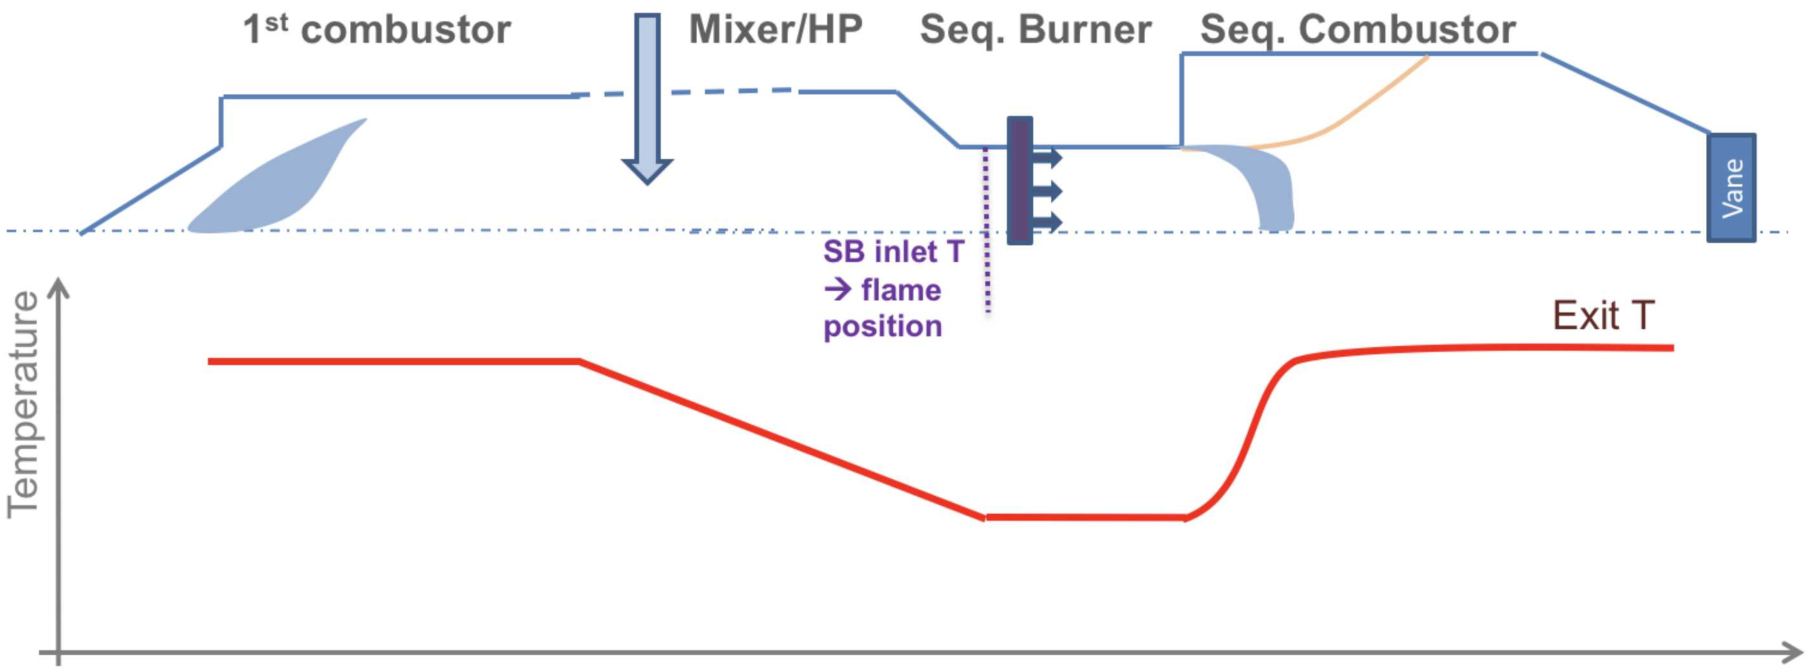

- Adjusting firing temperature of 1st stage allows control of t_{ign} in 2nd stage

Fuel and Load Flexibility Through Staging

Hydrogen fuel

- Flashback in 1st stage
- Early auto-ignition in 2nd stage

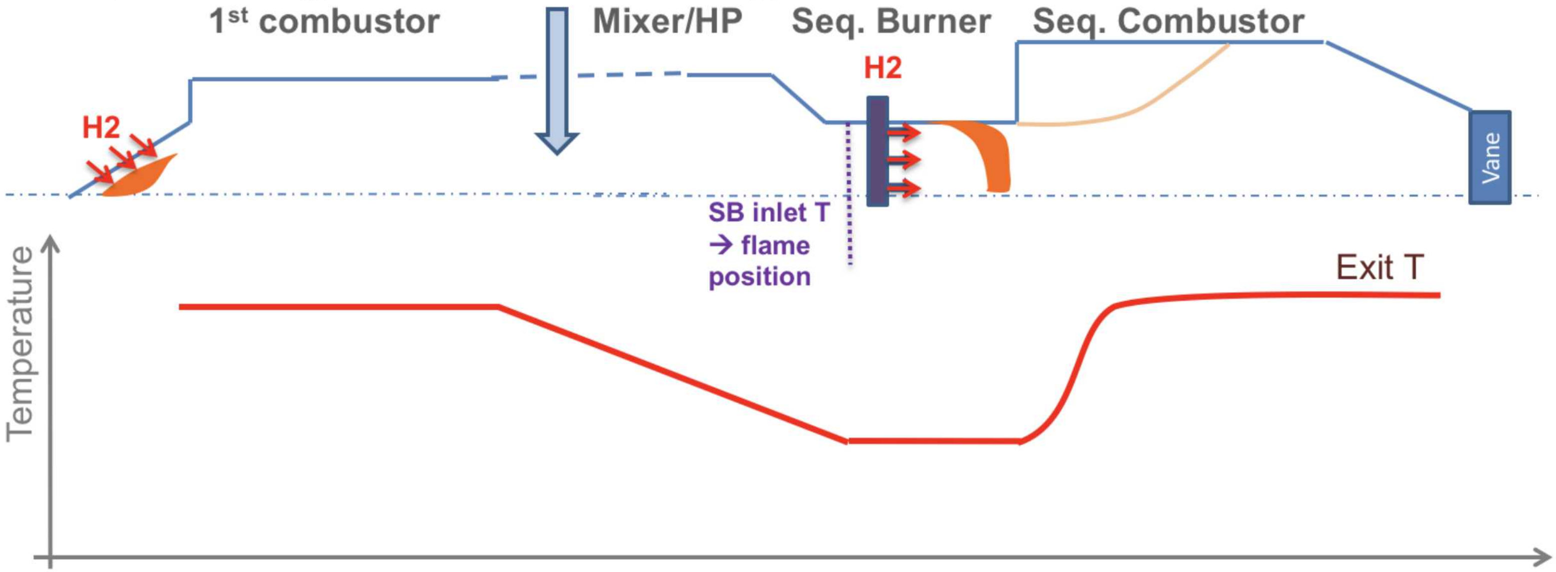

- 2nd stage is mainly auto-ignition stabilized
- 2nd stage inlet temperature needs to be decreased and not 2nd stage flame temperature

Fuel and Load Flexibility Through Staging

Hydrogen fuel

- Flashback in 1st stage
- Early auto-ignition in 2nd stage

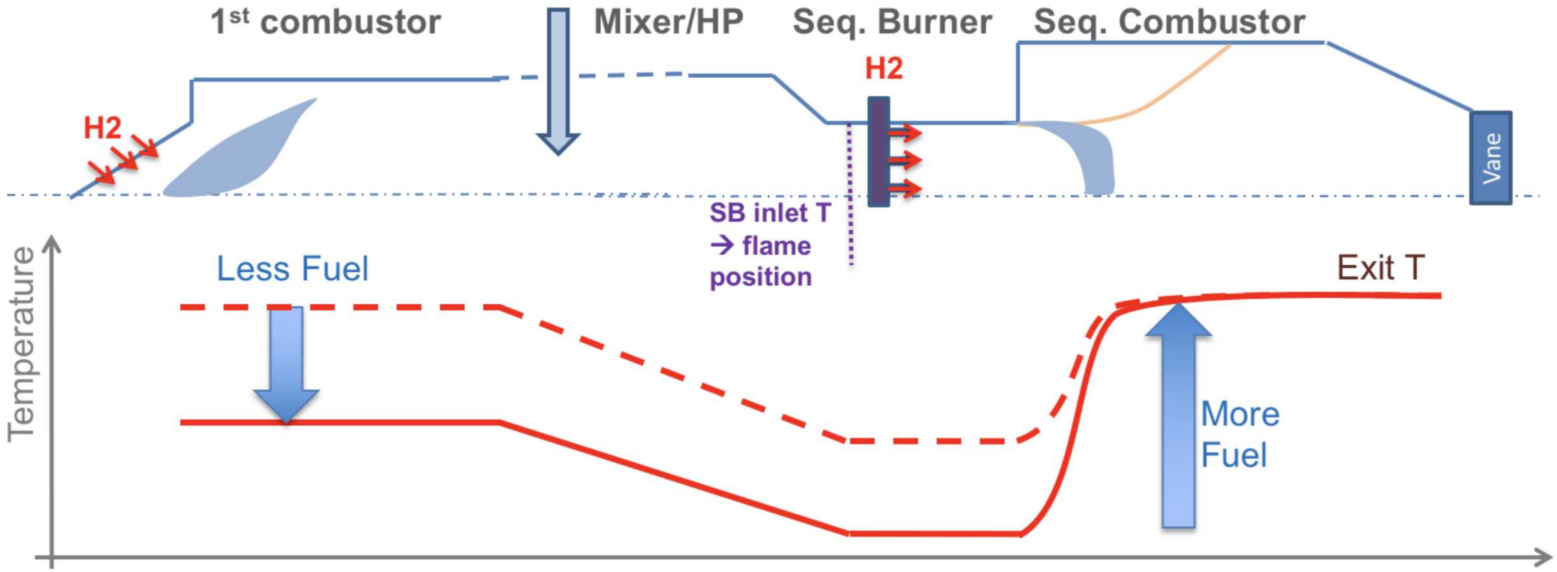

- 2nd stage is mainly auto-ignition stabilized
- 2nd stage inlet temperature needs to be decreased and not 2nd stage flame temperature
- 1st stage de-rating is compensated by shifting fuel to 2nd stage

Outline

- Motivation
- **Research paradigms**
- Governing equations, numerical methods, high performance computing

Focus: Turbulence-Chemistry Interactions

- Turbulence entrains, advects, strains and wrinkles a flame creating more area for burning

- Through turbulence cascade eventually reactants are molecularly mixed

- Chemical reactions are enhanced with mixing to a limit – extinction - and create heat release at dissipation scales

- Heat release, dilatation reduce turbulence intensity through density, and property changes and may alter spectral energy transfer

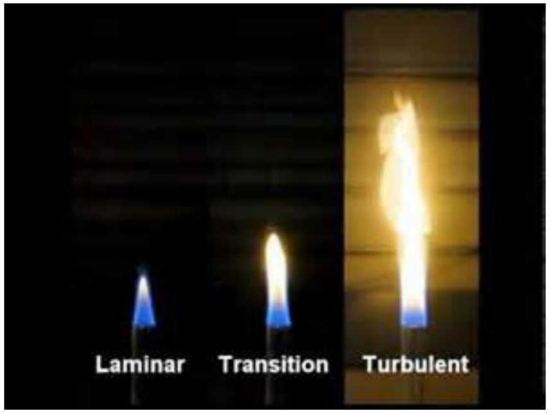

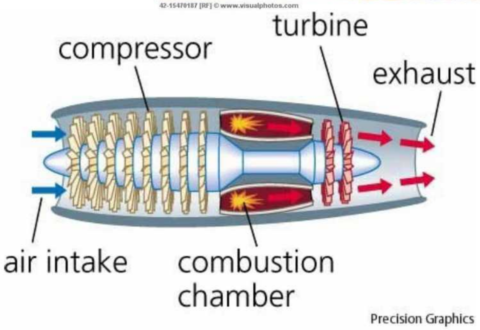

Precision Graphics

Combustion Research Paradigms: Experiment/Computation/Theory

- Experiments in laboratory scale flames/ignition at ambient and elevated pressure with advanced laser diagnostics are critical.

http://www.sandia.gov/ecn/

- High-fidelity simulations (DNS and LES) play a complementary role:
  - Precise control of parameters
  - Desired level of detail
  - Full specification of boundary and initial conditions (scalars, velocity)
  - Less expensive, faster design-cycle

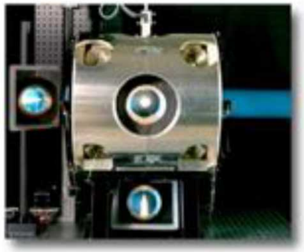

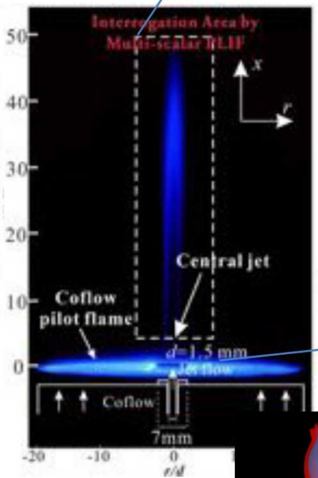

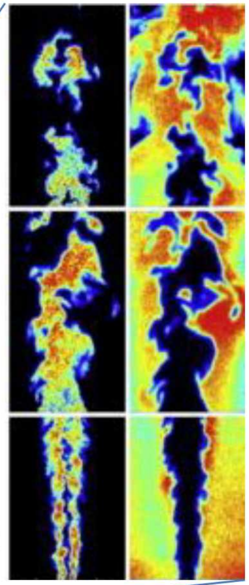

- Computational power and algorithms are advancing rapidly.

- Theory needed for interpretation, scaling, for identifying reduced manifolds and closure models

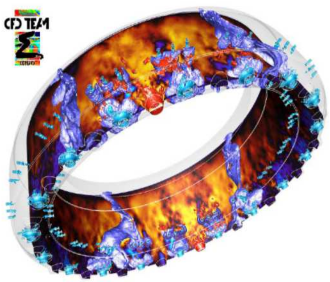

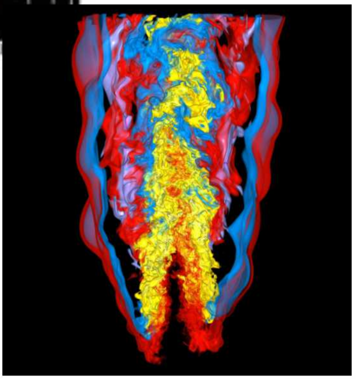

Turbulent Combustion: Computational Considerations

- Combustion in most practical devices occurs in turbulent flows where the ratio of inertial to viscous forces is large
- It is both:
  - Multi-scale (~ 4-6 decades in space and time).
  - Multi-physics (fluid dynamics, molecular transport, chemical kinetics, multi-phase).
- Governing equations represent a tightly coupled system with large dimensionality ($5 + N_{species}$).
- Computational cost becomes prohibitive very quickly.

Spatial Scales (Turbulence, Mixing, and Flames)

1. Kolmogorov lengthscales

$$\eta \approx \frac{\Lambda}{Re_t^{(3/4)}}; \quad \Lambda = k_1 L \quad L = N\Delta x$$

$$\eta > k_2 \Delta x \Rightarrow Re_t^{(3/4)} < \frac{k_1}{k_2} N \Rightarrow Re_t < \left(\frac{k_1}{k_2}\right)^{4/3} N^{4/9}$$

2. Batchelor lengthscales

$$\lambda_\beta = \frac{\eta}{\sqrt{Sc}} \quad Sc = \frac{\nu}{D} \approx \mathcal{O}(1)$$

Hydrogen-air, $Sc \approx 0.2$; n-heptane-air, $Sc \approx 2.4$

3. Chemical lengthscales:

$$\Delta x < \frac{\delta}{Q} \Rightarrow \frac{L}{\delta} < \frac{N}{Q} \quad Q \approx 20$$

Spatial Resolution – Nonreacting Turbulence

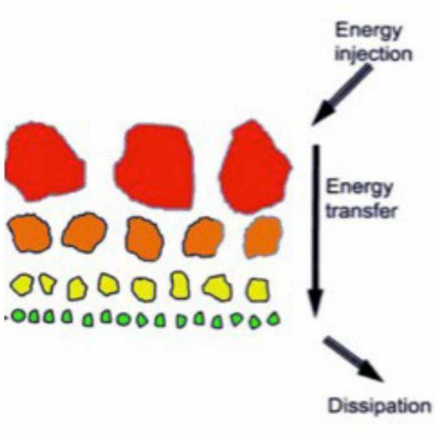

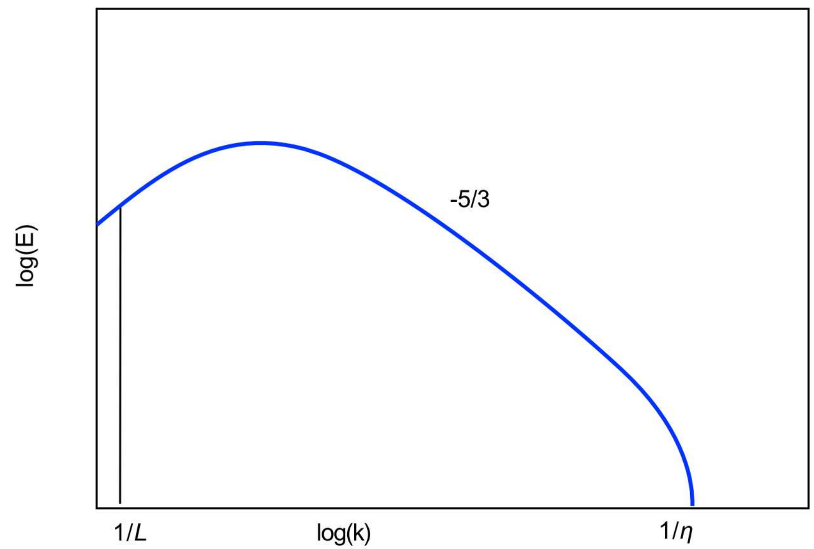

$$\Delta \approx \eta$$

$$N_{grid} \approx (L/\eta)^3$$

$$L/\eta \approx Re^{3/4}$$

$$\Rightarrow \quad \boxed{N_{grid} \approx Re^{9/4}}$$

Spatial Resolution Reacting Turbulence

- For turbulent combustion:
  - $\delta < \eta$

$$N_{grid} \approx Re^{9/4} (\eta/\delta)^3$$

  - No. of equations is large

$$N_{species} \approx O(10^1\text{-}10^2)$$

$$L \approx 10^{-1}, Re \approx 10^4, \eta \approx 10^{-4} m$$

$$\delta \approx (\text{diffusivity/flame speed}) \approx 10^{-5} m$$

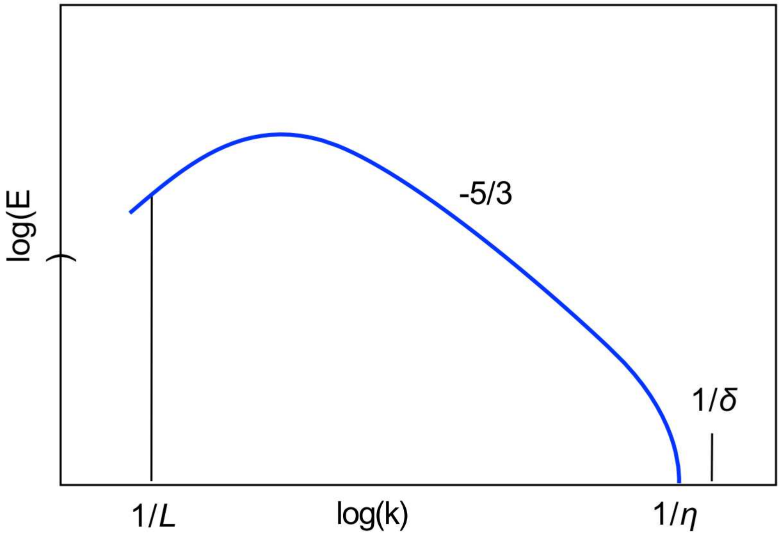

Temporal Resolution

1. Acoustic CFL

$$\Delta t < \frac{\Delta x}{a} \qquad \frac{\Delta t_a}{\Delta t_{u'}} = Ma$$

2. Advective CFL

$$\frac{\Delta t_a}{\Delta t_{u'}} = Ma$$

3. Chemical timescale
   - Flame timescale $\left(\tau_c \sim \frac{\delta}{s_L}\right)$
   - Species creation rates $\max\left(\dot{S}_i^{-1}\right)$
   - Reaction rates $\max\left(\dot{\omega}_j^{-1}\right)$
   - Eigenvalues of reaction rate jacobian

Multi-scale Modeling of Combustion Processes

- **Multi-scale modeling describes combustion processes, from quantum scales up to device-level, continuum scales**
- **Multi-scale Strategy:**
- **Use petascale computing power to perform direct simulation at the atomistic and fine-continuum scales (~4 decades), and develop new parameterizations that will enable bootstrapping information upscale**

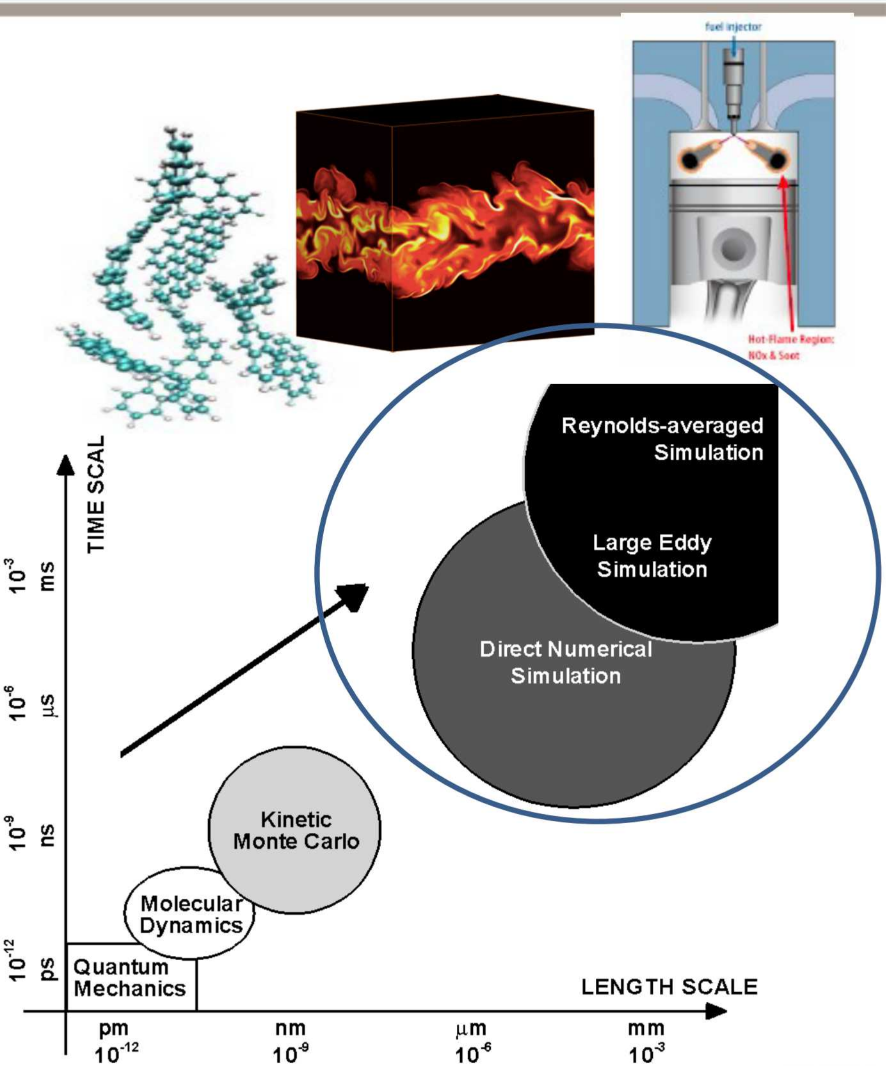

DNS – LES – RANS

- **Reynolds Averaged Navier Stokes (RANS)**
  - No attempt to resolve any scales
  - Inexpensive
- **Large Eddy Simulation (LES)**
  - Energy containing scales are resolved
  - Model subgrid physics
- **Direct Numerical Simulation (DNS)**
  - All continuum scales are resolved
  - Most expensive

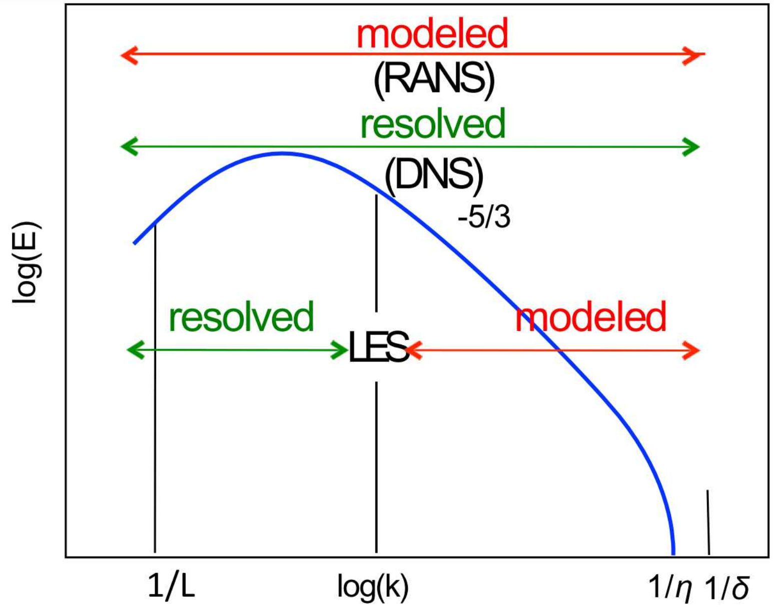

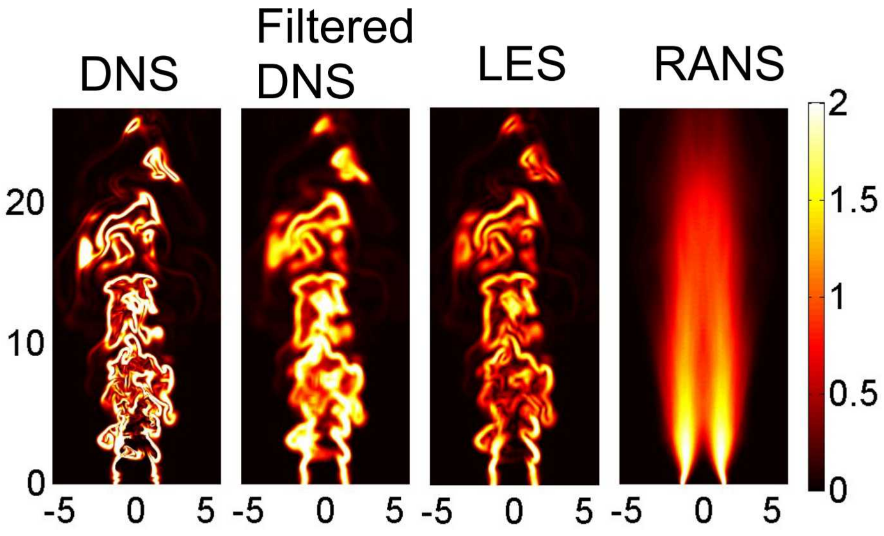

Combustion Simulation Approaches (DNS/LES/RANS)

- RANS and LES require models for

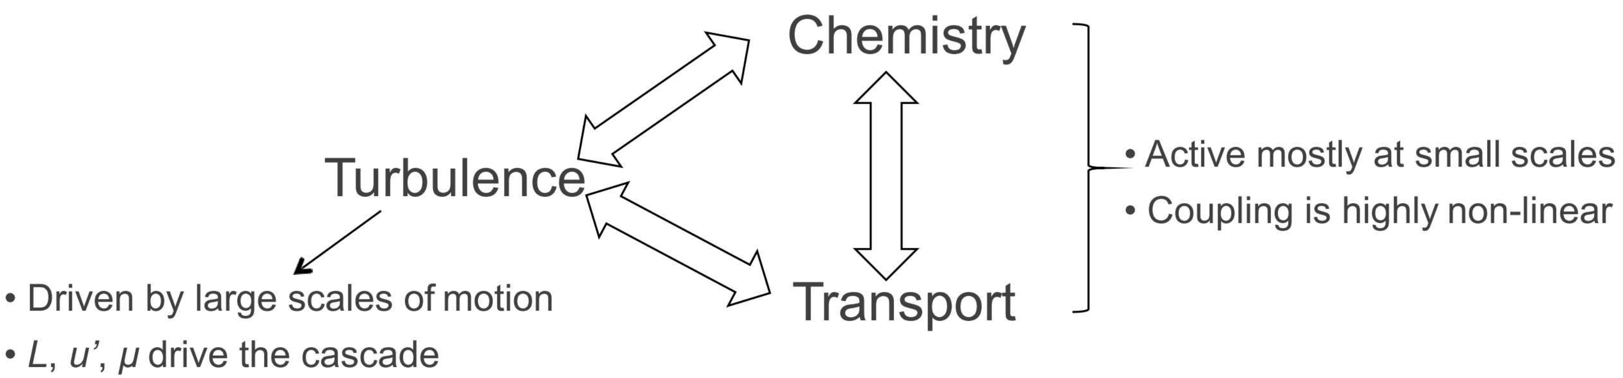

- DNS is often used to develop/validate models.
- DNS resolves all turbulence scales, no turbulence closure.
- DNS is not free of modeling (needs chemical and molecular transport models).

Outline

- Motivation
- Research paradigms
- **DNS mathematical formulation, numerical methods, and high performance computing**

Multi-scale Modeling of Combustion Processes

- Multi-scale modeling describes combustion processes, from quantum scales up to device-level, continuum scales

- Multi-scale Strategy:

  Use petascale computing power to perform direct simulation at the atomistic and fine-continuum scales (~4 decades), and develop new parameterizations that will enable bootstrapping information upscale

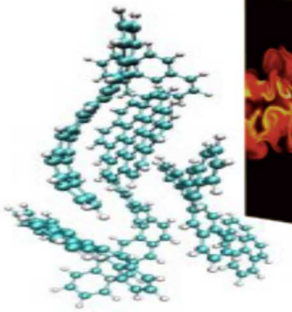

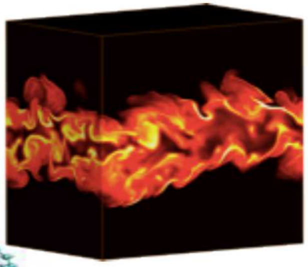

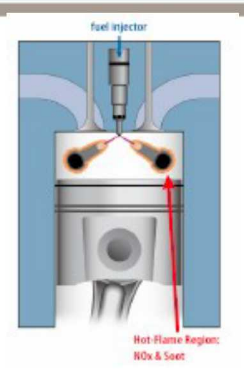

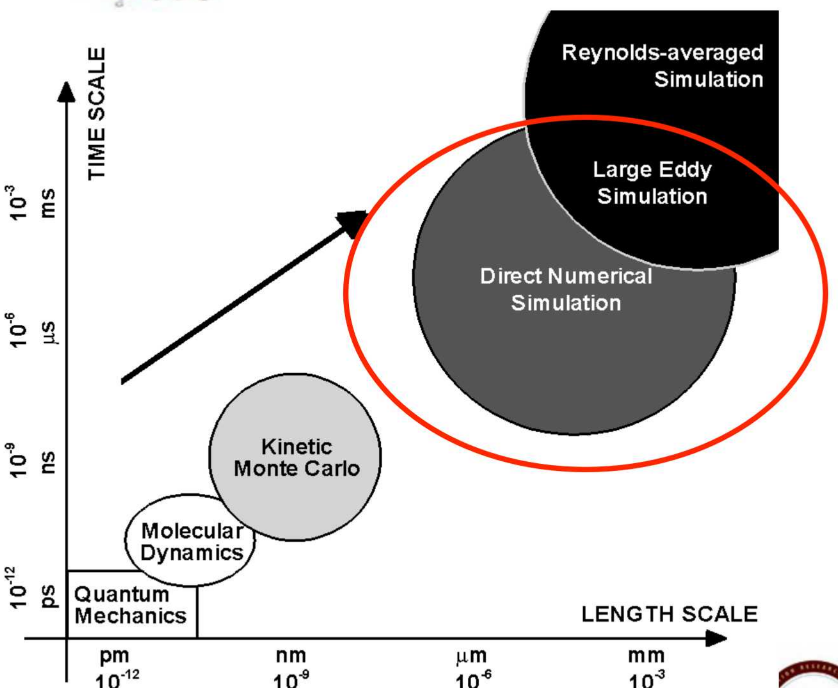

Governing Equations – Compressible

The equations governing reacting flows may be written in conservative form as

$$\frac{\partial \rho}{\partial t} = -\nabla_\beta \cdot (\rho \mathbf{u}_\beta), \tag{1}$$

$$\frac{\partial (\rho \mathbf{u}_\alpha)}{\partial t} = -\nabla_\beta \cdot (\rho \mathbf{u}_\alpha \mathbf{u}_\beta) + \nabla_\beta \cdot \boldsymbol{\tau}_{\beta\alpha} - \nabla_\alpha p + \rho \sum_{i=1}^{N_s} Y_i \mathbf{f}_{\alpha i}, \tag{2}$$

$$\frac{\partial (\rho e_0)}{\partial t} = -\nabla_\beta \cdot \left[\mathbf{u}_\beta (\rho e_0 + p)\right] + \nabla_\beta \cdot (\tau_{\beta\alpha} \cdot \mathbf{u}_\alpha) - \nabla_\beta \cdot \mathbf{q}_\beta + \rho \sum_{i=1}^{N_s} Y_i \mathbf{f}_{\alpha i} \cdot (\mathbf{V}_{\alpha i} + \mathbf{u}_\alpha), \tag{3}$$

$$\frac{\partial (\rho Y_i)}{\partial t} = -\nabla_\beta \cdot (\rho Y_i \mathbf{u}_\beta) - \nabla_\beta \cdot (\rho Y_i \mathbf{V}_{\beta i}) + W_i \dot{\omega}_i, \tag{4}$$

where ∇_β is the gradient operator in direction β, Y_i is the mass fraction of species i, W_i is the molecular weight of species i, $\tau_{\beta\alpha}$ is the stress tensor, $\mathbf{f}_{\alpha i}$ is the body force on species i in direction α, $\mathbf{q}_\beta$ is the heat flux vector, $\mathbf{V}_{\beta j}$ is the species mass diffusion velocity, $\dot{\omega}_i$ is the molar production rate of species i and e_0 is the specific total energy (internal energy plus kinetic energy),

$$e_0 = \frac{\mathbf{u}_\alpha \cdot \mathbf{u}_\alpha}{2} - \frac{p}{\rho} + h \tag{5}$$

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

Governing Equations - Compressible

$$\sum_{i=1}^{N_s} Y_i = 1. \tag{6}$$

Assuming an ideal gas mixture, the equation of state is given as

$$p = \frac{\rho R_u T}{W}, \tag{7}$$

where R_u is the universal gas constant and W is the mixture molecular weight given by

$$W = \left(\sum_{i=1}^{N_s} Y_i / W_i \right)^{-1} = \sum_{i=1}^{N_s} X_i W_i. \tag{8}$$

The species mass fractions (Y_i) and mole fractions (X_i) are related by

$$\frac{Y_i}{X_i} = \frac{W_i}{W}. \tag{9}$$

Relevant thermodynamic relationships between enthalpy and temperature for an ideal gas mixture include

$$h = \sum_{i=1}^{N_s} Y_i h_i, \quad h_i = h_i^0 + \int_{T_0}^{T} c_{p,i}\, \mathrm{d}T,$$

$$c_p = \sum_{i=1}^{N_s} c_{p,i} Y_i, \quad c_p - c_v = R_u / W,$$

Constitutive Relations

2.2. Constitutive relationships

The stress tensor, species diffusion velocities and heat flux vector in equations (2)–(4) are given by [2–4]

$$\tau_{\beta\alpha} = \tau_{\alpha\beta} = \mu\left[\nabla_\alpha \mathbf{u}_\beta + \nabla_\beta \mathbf{u}_\alpha\right] - \delta_{\alpha\beta}\left(\tfrac{2}{3}\mu - \kappa\right)\nabla_\gamma \cdot \mathbf{u}_\gamma, \qquad (10)$$

$$\mathbf{V}_{\alpha i} = \frac{1}{X_i}\sum_{j=1}^{N_s}\frac{Y_j}{X_j}D_{ij}\mathbf{d}_{\alpha j} - \frac{D_i^T}{\rho Y_i}\nabla_\alpha(\ln T), \qquad (11)$$

$$\mathbf{q}_\alpha = -\lambda\nabla_\alpha T + \sum_{i=1}^{N_s} h_i \mathbf{J}_{\alpha i} - \sum_{i=1}^{N_s}\frac{p}{\rho Y_i}D_i^T \mathbf{d}_{\alpha j}, \qquad (12)$$

where μ is the mixture viscosity, κ is the bulk viscosity, D_{ij} are the *multicomponent* diffusion coefficients, D_i^T is the thermal diffusion coefficient for species i, λ is the thermal conductivity, $\mathbf{J}_{\alpha i} = \rho Y_i \mathbf{V}_{\alpha i}$ is the species diffusive flux and $\mathbf{d}_{\alpha i}$ is the diffusion driving force for species i in direction α, given by [2–4]

$$\mathbf{d}_{\alpha i} = \underbrace{\nabla_\alpha X_i}_{1} + \underbrace{(X_i - Y_i)\nabla_\alpha(\ln p)}_{2} + \underbrace{\frac{\rho Y_i}{p}\left[\mathbf{f}_{\alpha i} - \sum_{j=1}^{N_s} Y_j \mathbf{f}_{\alpha j}\right]}_{3}. \qquad (13)$$

The driving force vector involves thermodynamic forces generated by gradients in concentration (term 1), gradients in pressure (term 2) also called 'barodiffusion', and due to any body force such as an electrical or gravitational field (term 3). Equation (13) allows for the possibility that the force on each species, $\mathbf{f}_{\alpha i}$, is different, though in the case of a gravitational field, $\mathbf{f}_{\alpha j} = \mathbf{g}_\alpha$, and term 3 is identically zero. In the following sections, we will consider the fluxes given in equations (10)–(12) in more detail.

Mass Diffusion Flux

2.4. Mass diffusion flux

It should be noted that all diagonal components of the multicomponent diffusion coefficient matrix (D_{ii}) are identically zero [2]. Also, the diffusive fluxes and driving forces for all species must sum to zero,

$$\sum_{i=1}^{N_s} \mathbf{J}_{\alpha i} = \sum_{i=1}^{N_s} \rho Y_i \mathbf{V}_{\alpha i} = 0, \qquad \sum_{i=1}^{N_s} \mathbf{d}_{\alpha i} = 0. \tag{15}$$

Equation (11) is often approximated as [2–5]

$$\mathbf{V}_{\alpha i} = -\frac{D_i^{\mathrm{mix}}}{X_i}\mathbf{d}_{\alpha i} - \frac{D_i^T}{\rho Y_i}\nabla_\alpha \ln T, \tag{16}$$

where D_i^{mix} are 'mixture-averaged' diffusion coefficients given in terms of the binary diffusion coefficients ($\mathcal{D}_{ij}$) and the mixture composition as

$$D_i^{\mathrm{mix}} = \frac{1 - X_i}{\sum_{j \neq i} X_j / \mathcal{D}_{ij}}, \tag{17}$$

where the binary coefficient matrix is symmetric ($\mathcal{D}_{ij} = \mathcal{D}_{ji}$), and the diagonal elements are zero ($\mathcal{D}_{ij} = 0$). If we assume that body forces act in the same manner on all species and baro-diffusion (term 2 in (13) is negligible, then (13) becomes $\mathbf{d}_{\alpha i} = \nabla_\alpha X_i$. If we further neglect the Soret effect, (the second term in (11) and (16), then (16) reduces to

$$\mathbf{V}_{\alpha i} = -\frac{D_i^{\mathrm{mix}}}{X_i}\nabla_\alpha X_i, \tag{18}$$

Mass Diffusion Flux / Heat Flux

which, using (9), can be expressed in terms of mass fractions as

$$\mathbf{V}_{\alpha i} = -\frac{D_i^{\mathrm{mix}}}{Y_i}\left[\nabla_\alpha Y_i + \frac{Y_i}{W}\nabla_\alpha W\right]$$
$$= -\frac{D_i^{\mathrm{mix}}}{Y_i}\left[\nabla_\alpha Y_i - Y_i W \sum_{j=1}^{N_s} \frac{\nabla_\alpha Y_j}{W_j}\right]. \quad (19)$$

Studies on the effects of thermal diffusion suggest that the Soret effect is much more important for premixed flames than for non-premixed flames, the Dufour effect is of little importance in either premixed or non-premixed flames [6].

2.5. Heat flux

The heat flux vector is comprised of three components representing the diffusion of heat due to temperature gradients, the diffusion of heat due to mass diffusion and the Dufour effect [2, 4–7]. In most combustion simulations, the Dufour effect is neglected, and (12) may be written as

$$\mathbf{q}_\alpha = -\lambda \nabla_\alpha T + \sum_{i=1}^{N_s} h_i \mathbf{J}_{\alpha i}. \quad (20)$$

Direct Numerical Simulation Code – S3D

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

- **Used to perform first-principles-based DNS of reacting flows**
- **Solves compressible reacting Navier-Stokes equations**
- **High-fidelity numerical methods**
- **Detailed reaction kinetics and molecular transport models**
- **Multi-physics: sprays, radiation and soot**
- **Ported to all major platforms, scales well**
- **Particle tracking cabability**

DNS ⟷ ***Physical models*** ⟷ ***Engineering CFD codes (RANS, LES)***

S3D Numerical Methods

- Finite difference solver (physical space).
- Rectangular Cartesian regular mesh.
- Explicit 8th order central difference (Kennedy & Carpenter, *Appl. Num. Math.,* 1994, vol. 14).
- Explicit 4th order 6-stage Runge-Kutta (Kennedy & Carpenter, *Appl. Num. Math.,* 2000, vol. 35).
- 10th order filter to remove spurious numerical noise.
- Δx, Δt resolve all relevant spatio-temporal scales (fluid-dynamic, thermo-chemical, acoustic).
- Boundary Conditions:
  - NSCBC for Non-reflecting inflow/outflow (Poinsot and Lele, *J. Comp. Phys.* 1992, Sutherland and Kennedy, *J. Comp. Phys.* 2003, Yoo and Im, *Comb. Theory Mod.* 2007)
  - Periodic
  - Isothermal or adiabatic wall

Chemical Complexity vs Scale

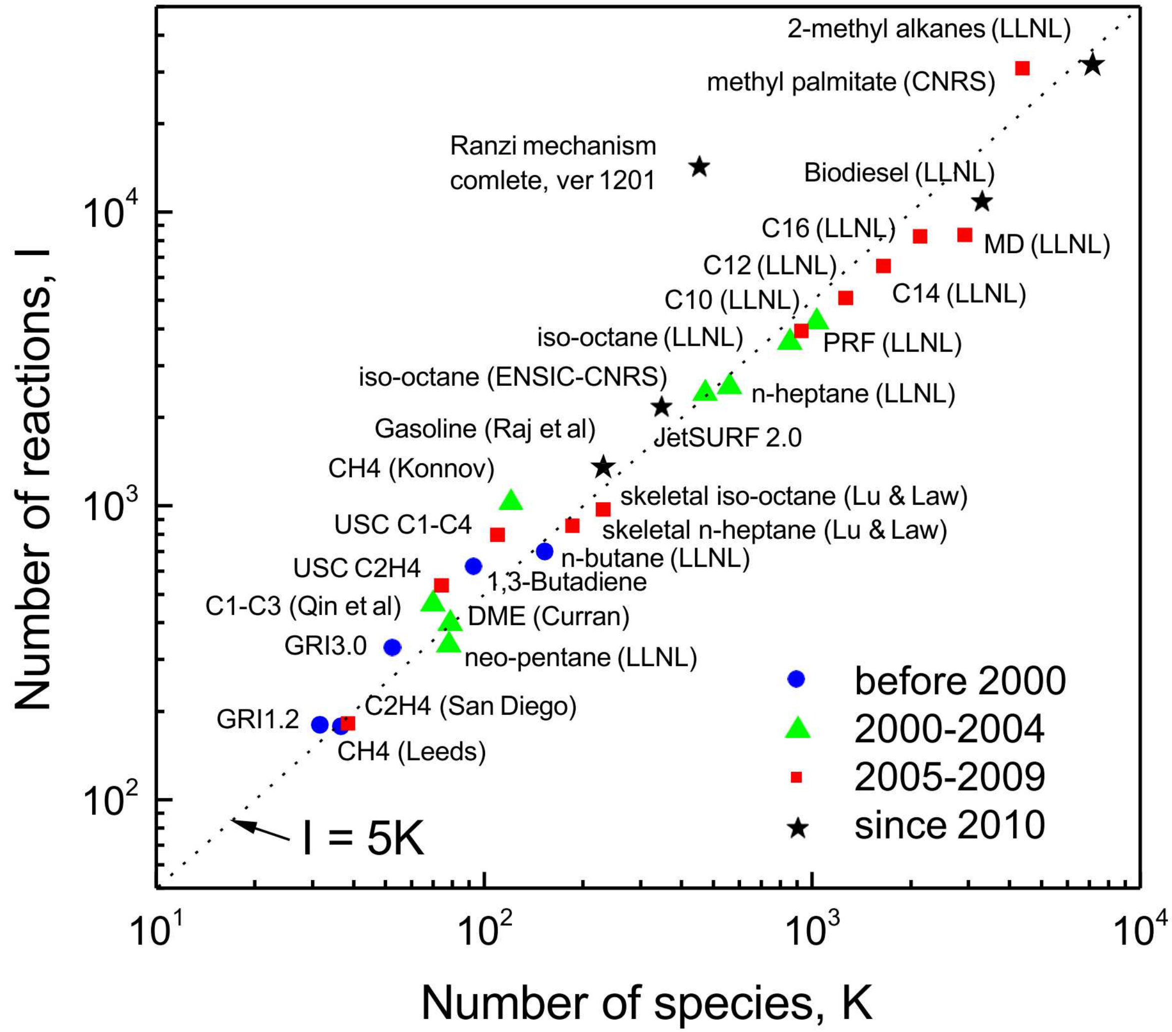

T. Lu and C. K. Law, *PECS 35,* 2009

A Systematic Procedure for Dimension Reduction & Stiffness Removal

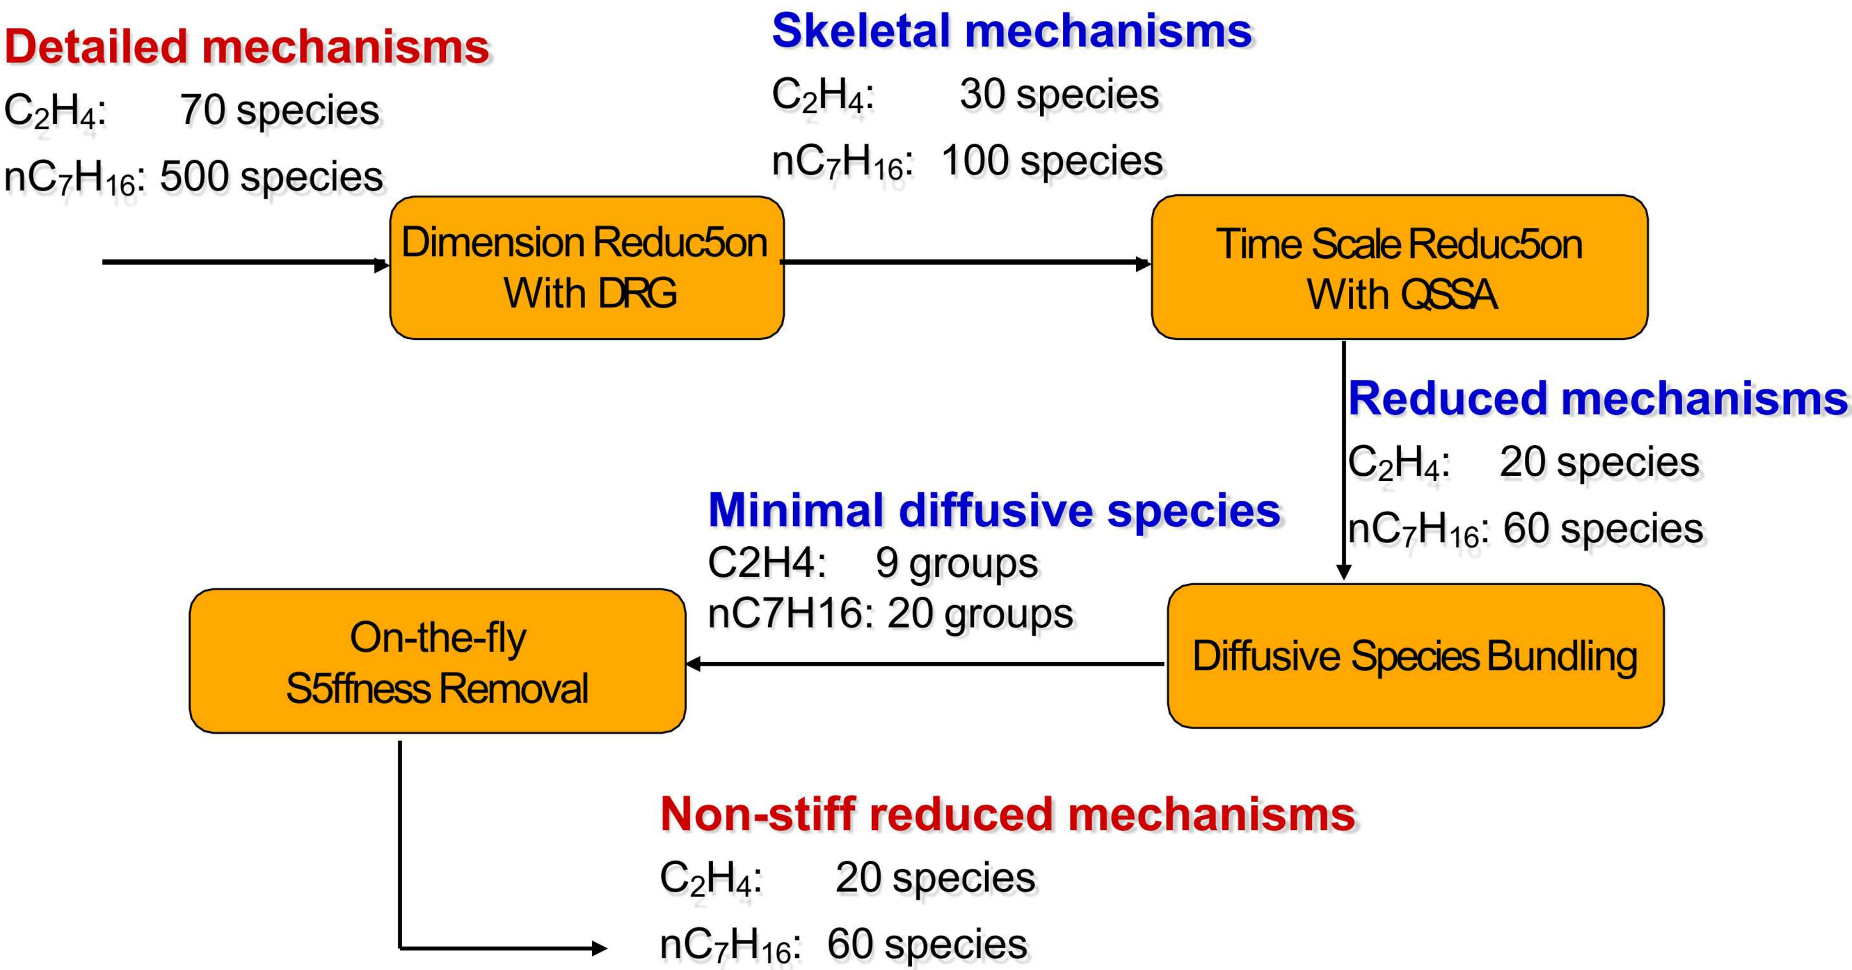

Lu and Law 2005, Lu and Law 2006,2008, Lu and Law 2007, Lu et al. 2009

S3D Parallelization – Data Parallelism

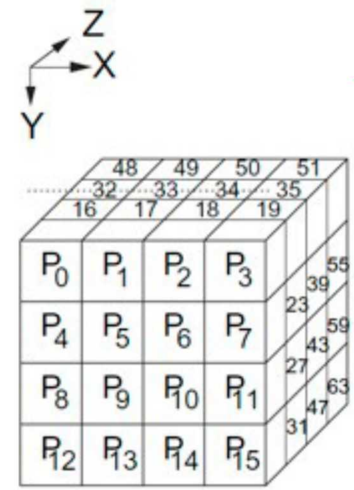

- Algorithm:
  - fixed mesh – trivial 3D domain decomposition (each MPI process computes on a piece of the 3D domain)
  - MPI only parallelism (hybrid MPI+OpenMP, Open ACC, OpenMP + CUDA).
  - ideal load balancing (all processes have same amount of work).
  - communication is almost entirely between nearest neighbors – ghost cells, large message sizes, nonblocking sends/receives
  - parallel collective I/O
  - per node performance limited by memory bandwidth.
  - extremely scalable and portable (CRAY, IBM, INTEL).

On JAGUAR, at NCCS (Oak Ridge National Lab)

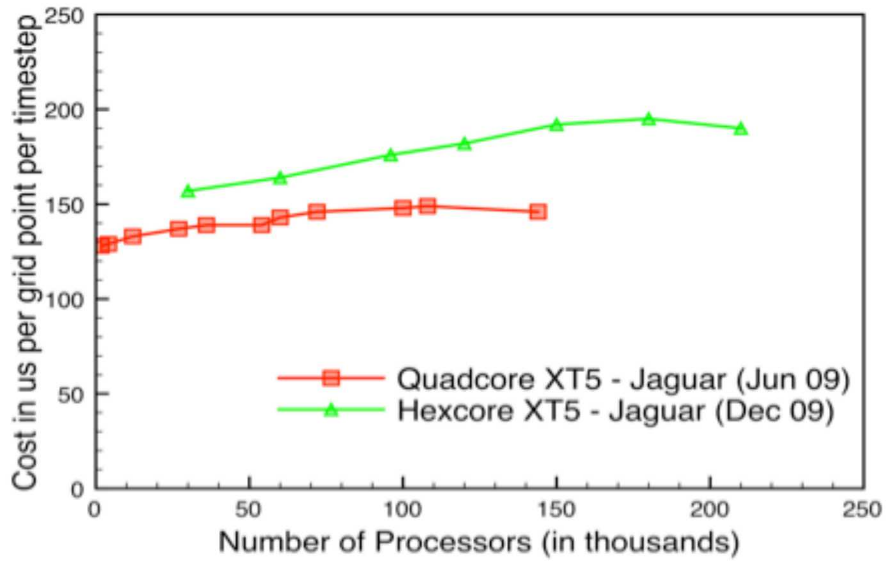

Domain Decomposition and Bulk Synchrony

Bulk Sychronous

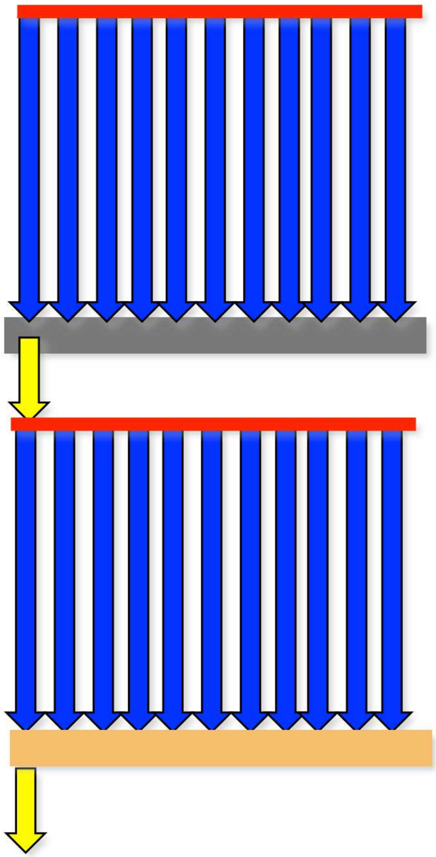

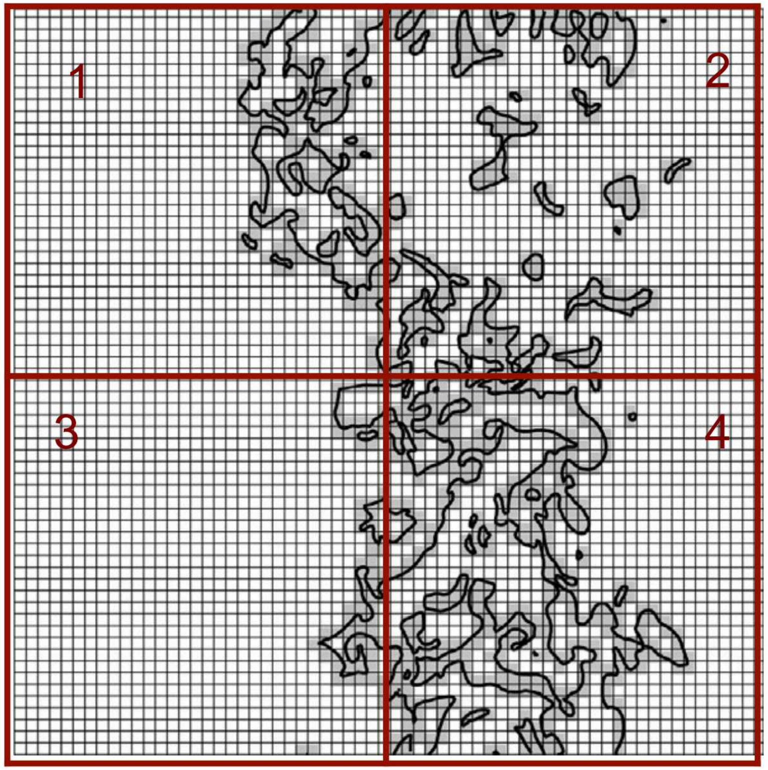

.Each MPI process is in charge of a piece of the 3D domain.
Inter-processor communication is only between nearest neighbors

Compressible DNS Solver – Main Loop

- Computes rate of change of N conserved quantities at every grid point
  - d/dt (Q_k) = (Advection) + (Diffusion) + (Source)
  - Sum of all the terms that contribute to the time derivative is called the RHS
- d/dt (Q_k) is integrated explicitly in time through Runge-Kutta (4^{th} order – 6 stage, Kennedy and Carpenter)
- RHS contains multiple terms that are functions of Q_k, variables derived from Q_k
- Advection and diffusion require spatial finite differencing and nearest neighbor ghost zone communication, MPI
- Source terms are point-wise functions
- Thermodynamic, chemical and molecular transport properties are point-wise functions of Q_k

Computational Intensity of Chemistry

- Called as ckwyp or getrates (Chemkin)
- Chemical reaction rate computed using Arrhenius model
  - A + B = C+ D
  - Forward reaction rate = $C*[A]*[B]*T^{a}exp(-Ta/T)$
  - Equilibrium constant gives reverse reaction rates
  - More terms for third body efficiency, collision efficiency, pressure corrections …
- The source term for a species is the sum of the rates of all reactions in which it participates
- The kernel uses exp/log heavily

Mechanism	Reactions	Species	Unique	QSSA	Stiff
H2	15	9	9	-	-
DME	175	39	30	9	22
Heptane	283	68	52	16	27
PRF	861	171	116	55	93

Figure 1: Summary of S3D Chemical Mechanisms

Computational intensity of DNS scales with Moore's Law

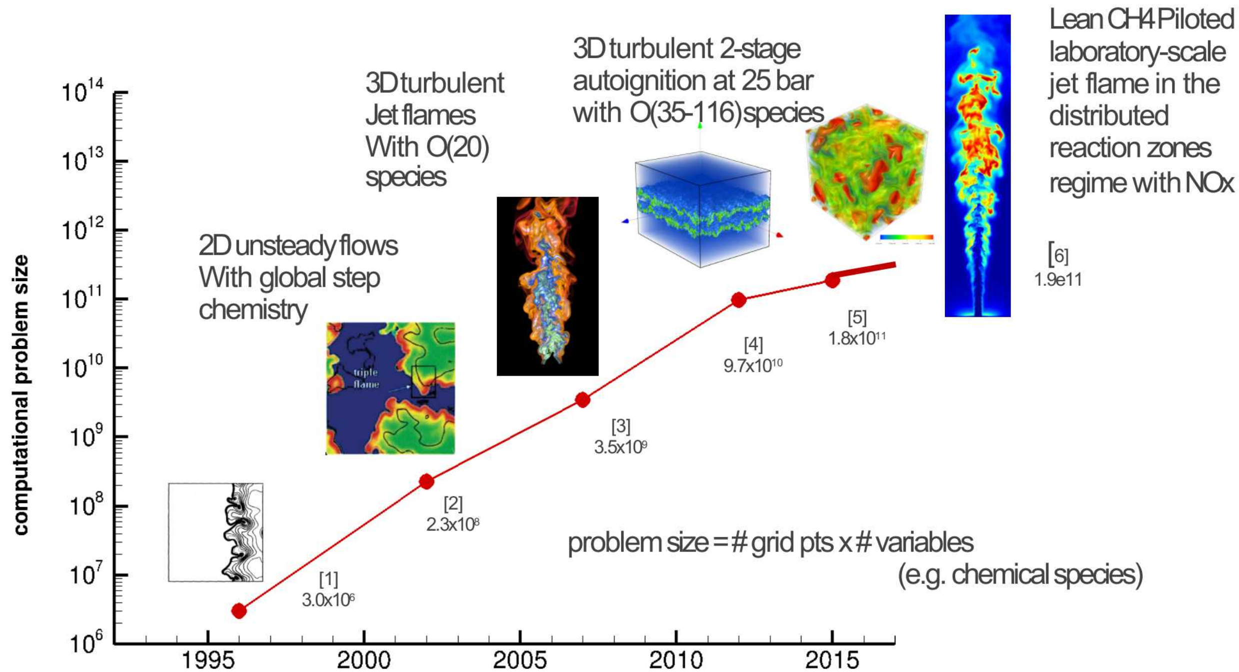

1 T.Echekki, J.H. Chen, *Comb. Flame*, 1996, vol.106.
2 T.Echekki, J.H, Chen, *Proc. Comb. Inst.*, 2002, vol. 29.
3 R. Sankaran, E.R. Hawkes, J.H. Chen, *Proc. Comb. Inst.*, 2007, vol. 31.
4 E.R. Hawkes, O. Chatakonda, H. Kolla, A.R. Kerstein, J.H. Chen, *Comb. Flame*, 2012, vol. 159.
5 2015 submission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

DNS Benchmarks (Physical Insights & Model Development)

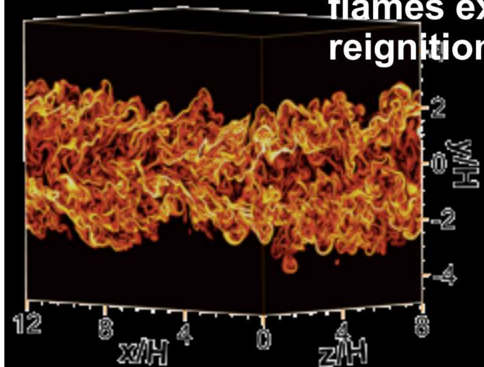

CO/H_2 and C2H4 jet flames extinction/ reignition

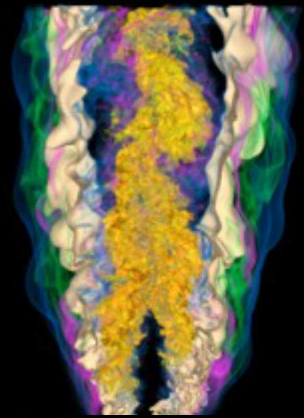

Lifted H2 and C2H4 jet flames in hot coflow

Di-Methyl Ether HCCI

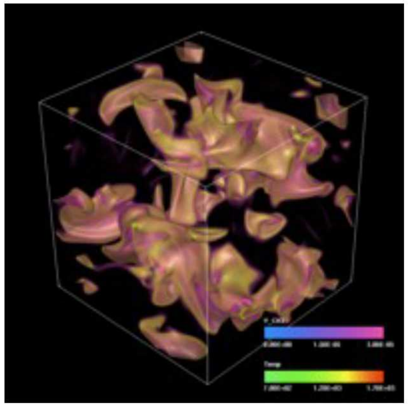

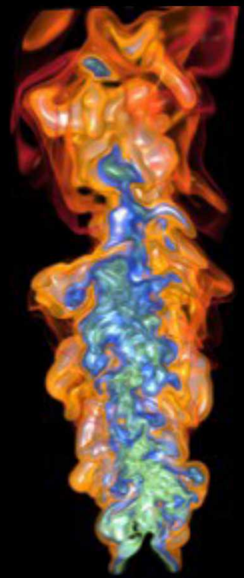

Lean premixed and stratified CH4/air Bunsen flames

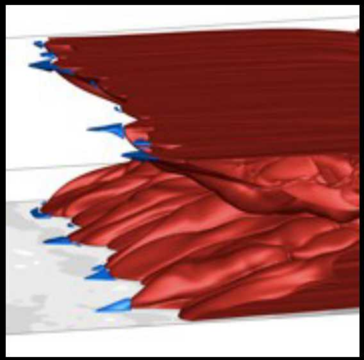

H2 Flame-Boundary Layer Interaction - Flashback

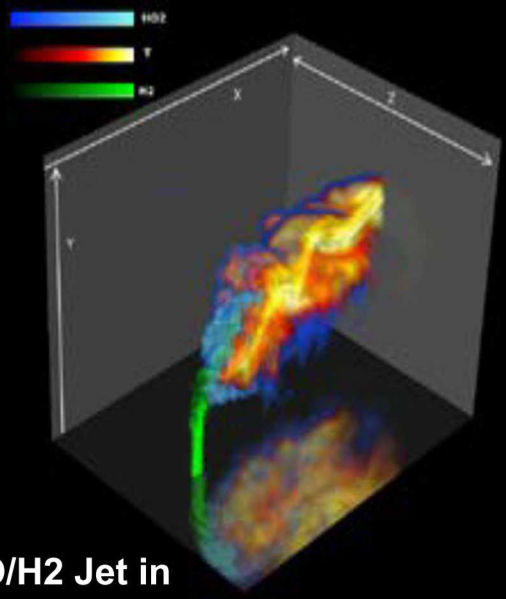

H_2 and CO/H2 Jet in Crossflow

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

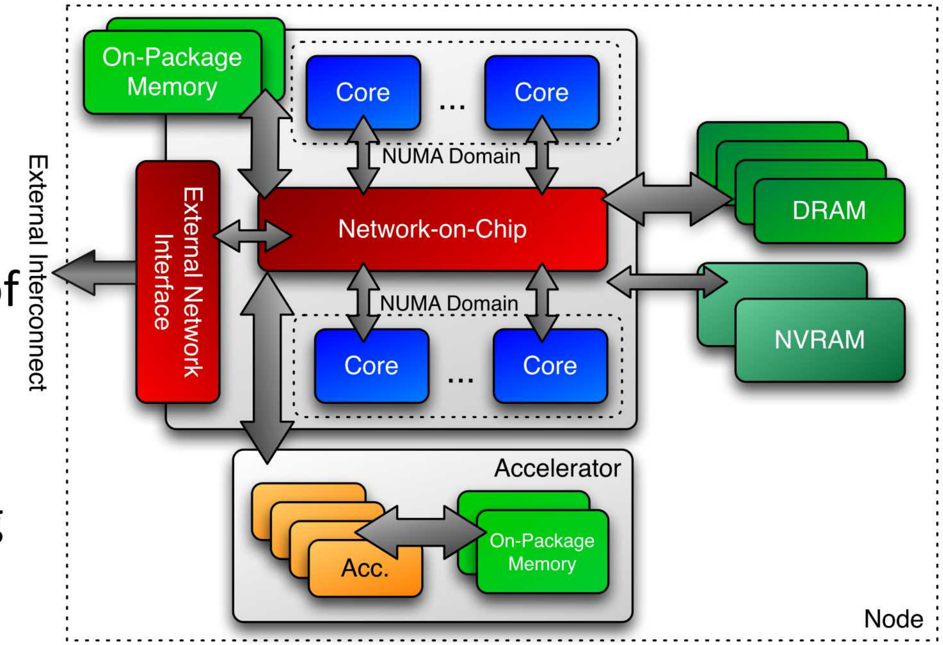

Conceptual model of future HPC node

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

DOE Exascale Computing Project (ECP) will achieve capable exascale machines 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

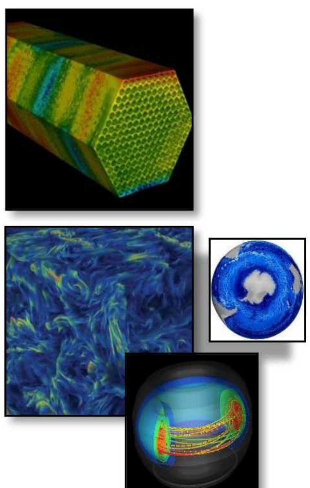

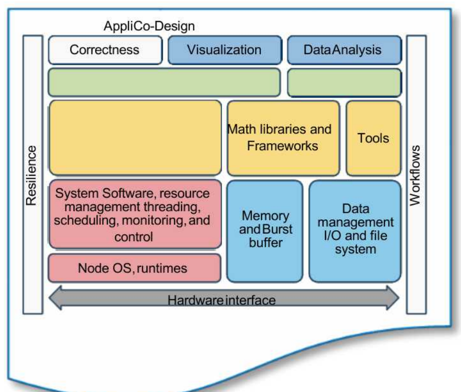

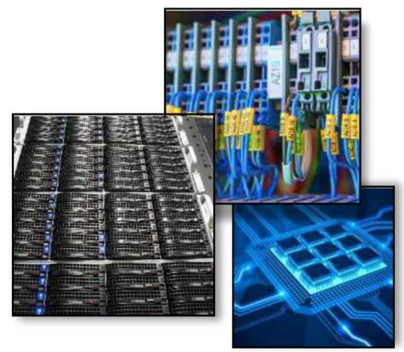

ECP's work encompasses applications, system software, hardware technologies and architectures, and workforce development

From Paul Messina's ASCAC talk April 19, 2017

ECP application: transforming combustion science and technology through exascale simulation (Pele)

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

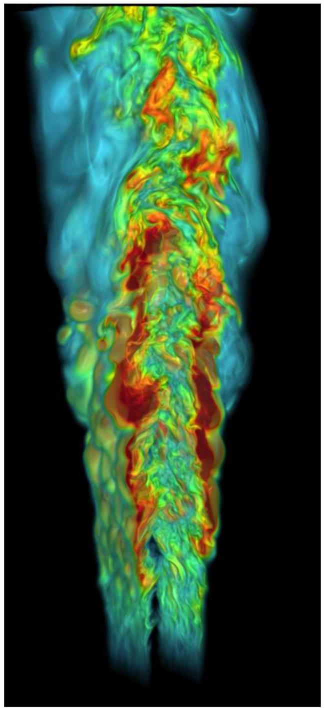

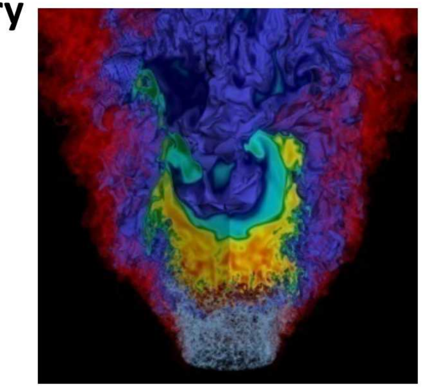

Effects of reactivity stratification at:

- high pressure
- high turbulence
- fuel blends

on:

- ignition delay
- combustion rates
- emissions

Automated Mechanism Generation

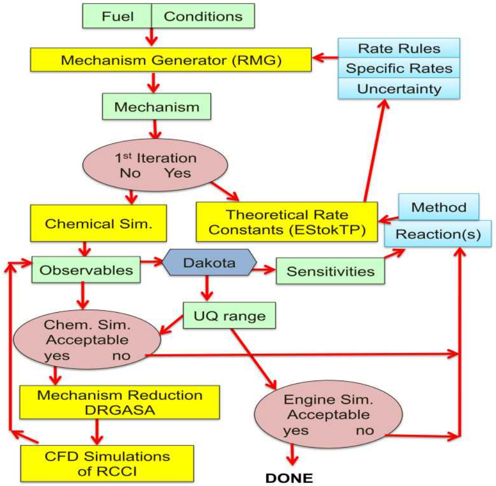

S3D: multi-block compressible reacting DNS multi-physics validation: spray, soot, radiation

Argonne NATIONAL LABORATORY · BERKELEY LAB Lawrence Berkeley National Laboratory · OAK RIDGE National Laboratory · Sandia National Laboratories · NREL NATIONAL RENEWABLE ENERGY LABORATORY

Programming Environment Critical to Performance

Effective use of exascale hardware will require programming environment that effectively maps algorithms to hardware

- Driven by programmability of combustion applications and characterization of algorithms on different designs of architectures
  - Simplify programming to express locality and independence
  - Automate discovery of parallelism and hide latencies
  - Simplify programming of extensible workflows, block-structured PDE's, analytics, UQ for performance, scalability, portability, and productivity on heterogeneous architectures

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

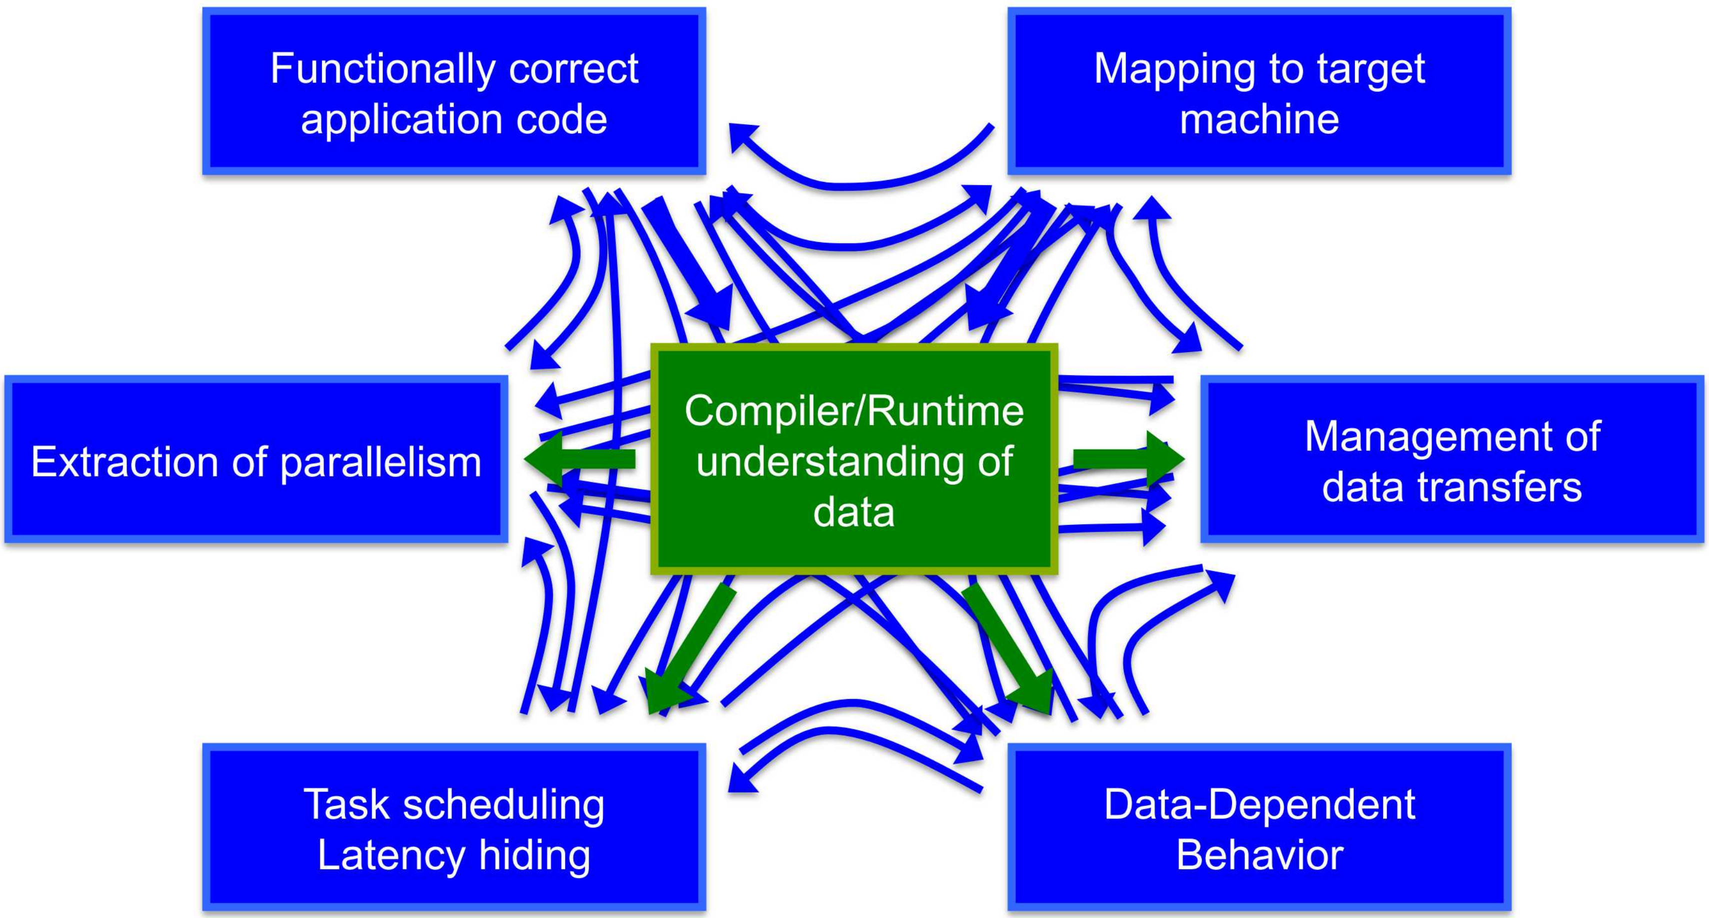

Legion Programming System applied to S3D

- 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

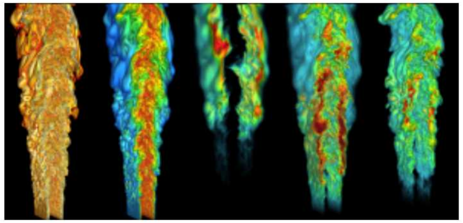

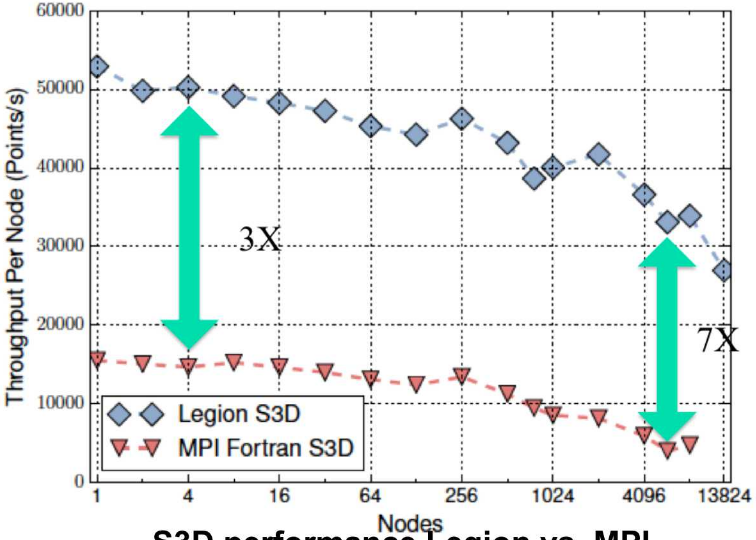

S3D performance Legion vs. MPI

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.

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

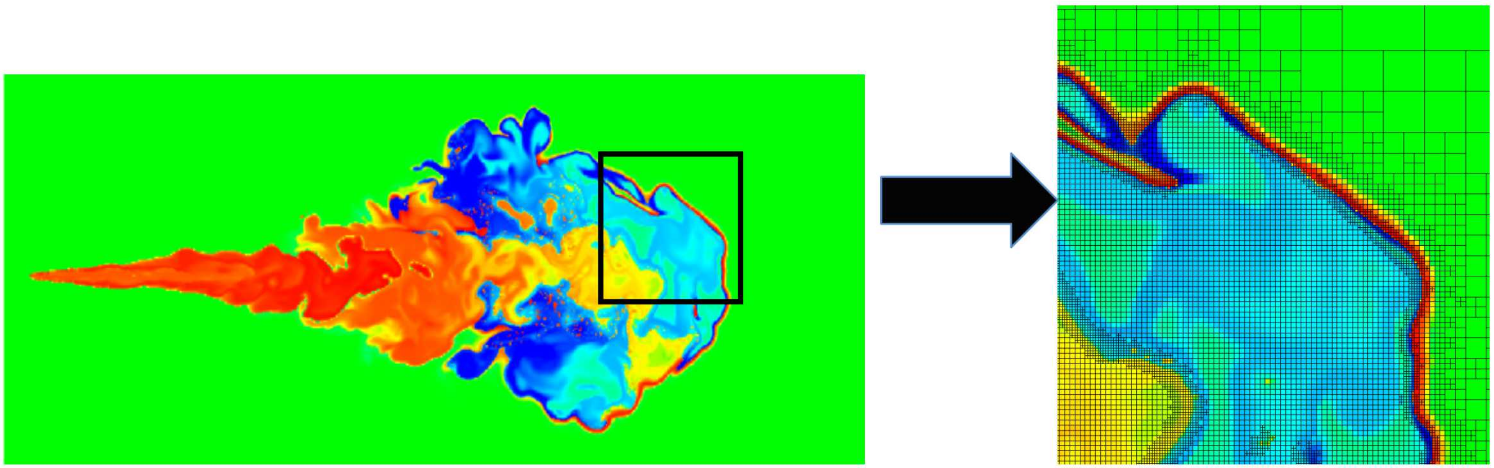

Run CEMA at each time step as a diagnostic to steer mesh refinement
CEMA computation takes longer than a single explicit RK stage (6 stages/timestep)
Dividing CEMA across RK stages and interleaving with other computation so as not to impact other critical operations would be hard to schedule manually
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 (Titan & Piz Daint)

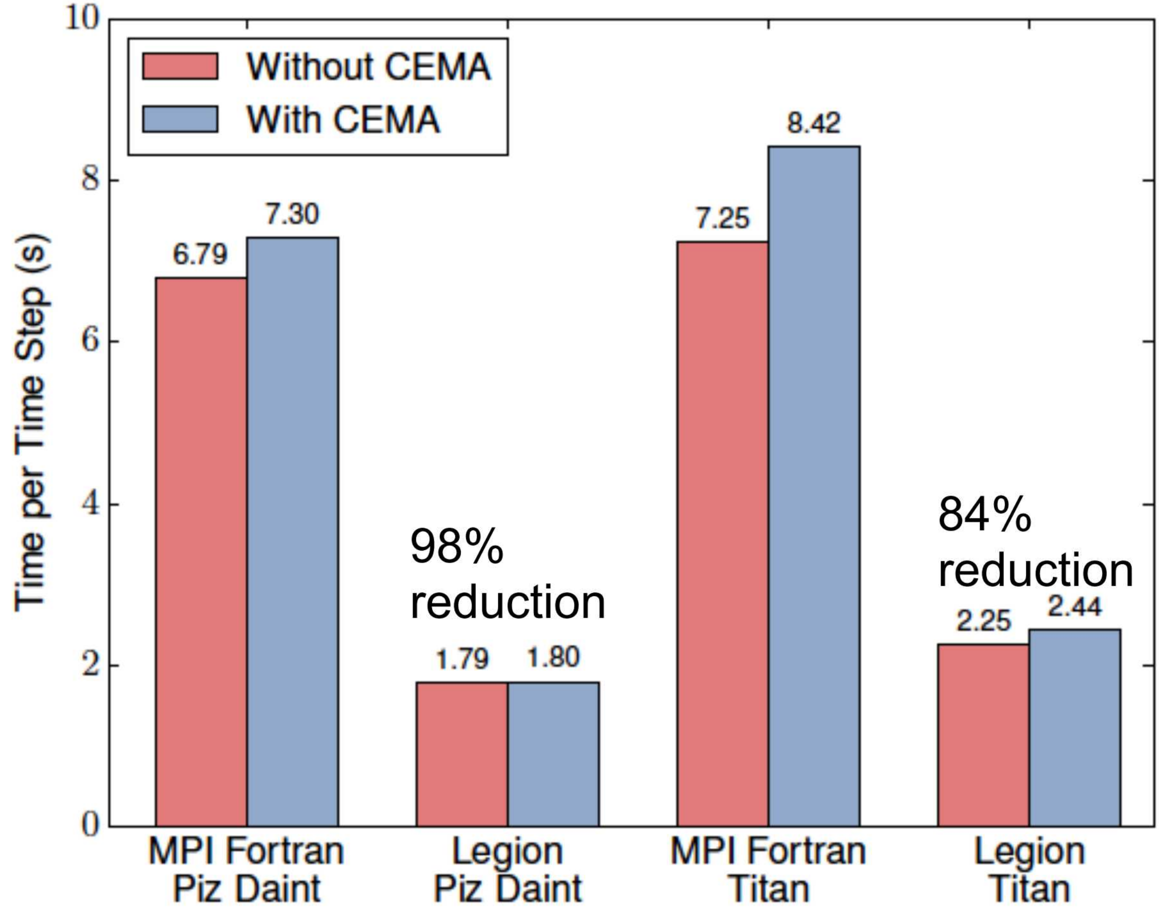

Legion S3D Lessons Learned

- 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
    - Productivity requires higher level abstraction layer for scientists to write in

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