SAND2008-4886C

Recent S3D Science and Performance

PERI SciDAC Meeting

Seattle, Washington

July 18, 2008

Ray Grout, Jaqueline H. Chen
Combustion Research Facility, SNL

Ramanan Sankaran
Center for Computational Sciences, ORNL

Sponsored by: DOE Office of Basic Energy Sciencesand Advanced Scientific Computing Research

Multi-scale Modeling

- New engine technologies, combustion modes
  – Sound scientific understanding essential to develop predictive, validated multi-scale models

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

- Terascale computing can directly resolve 3 decades of scale

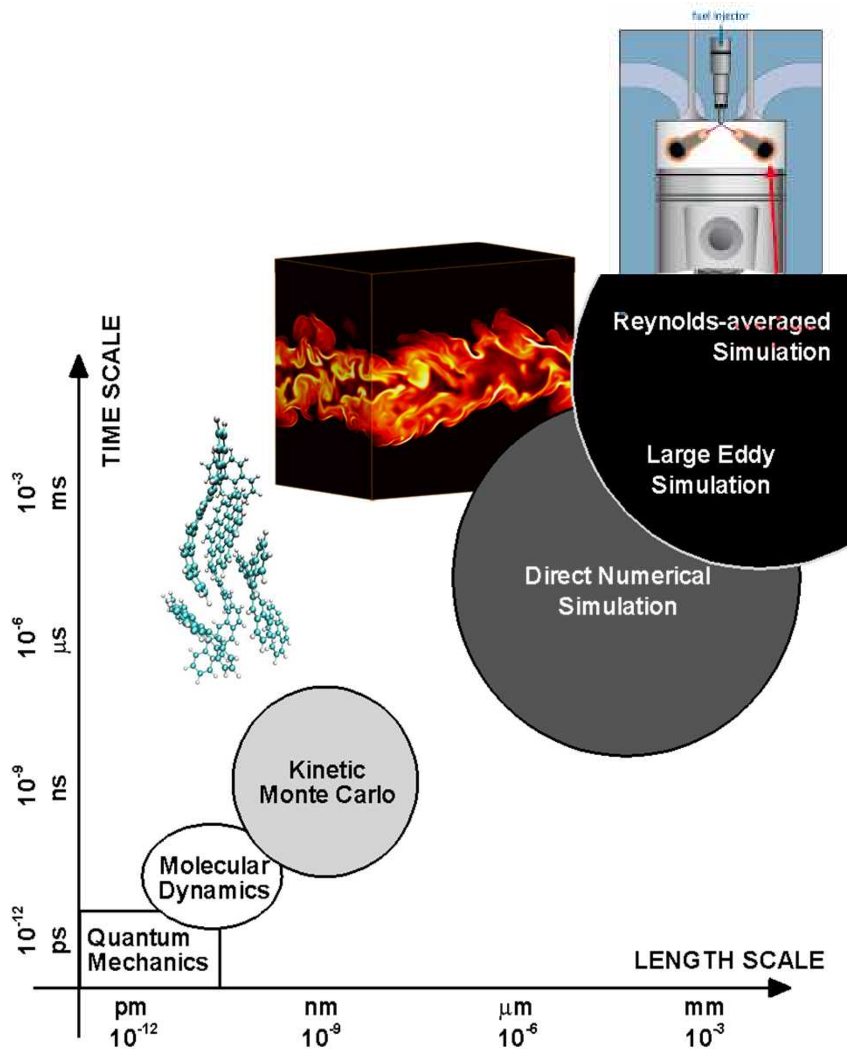

Sandia National Laboratories

Community Data for Model Validation

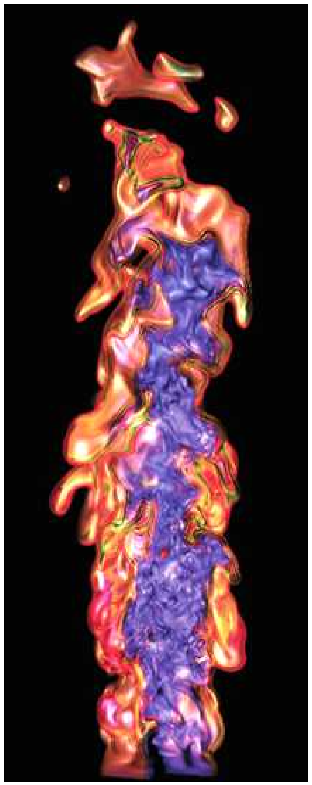

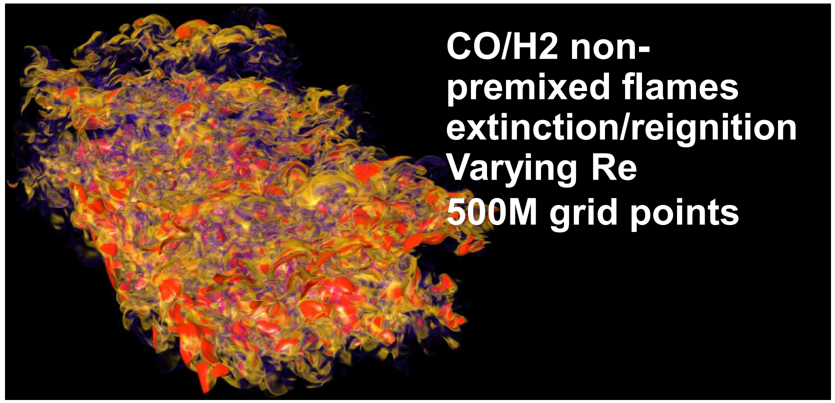

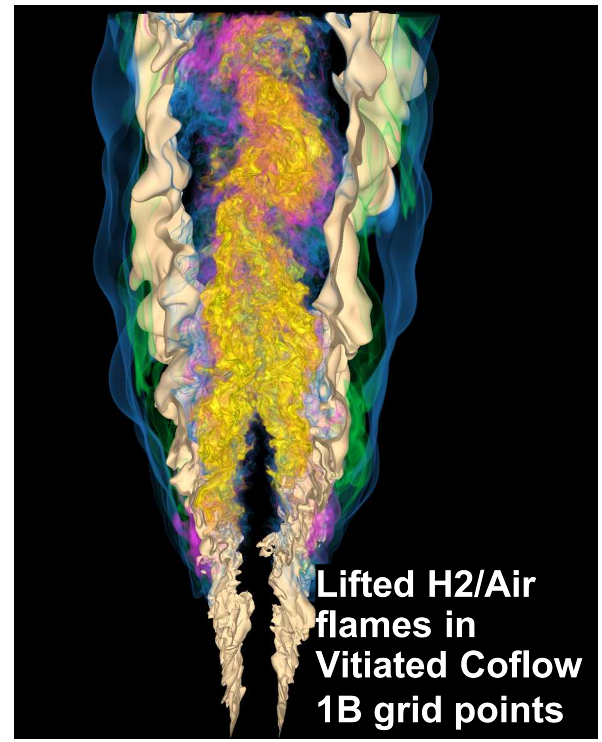

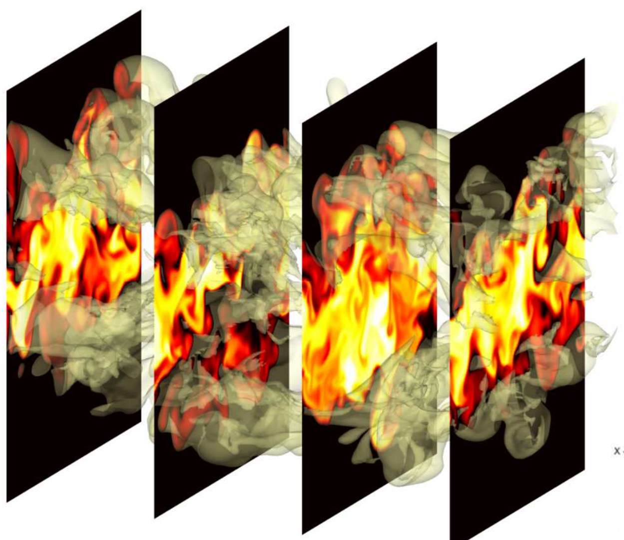

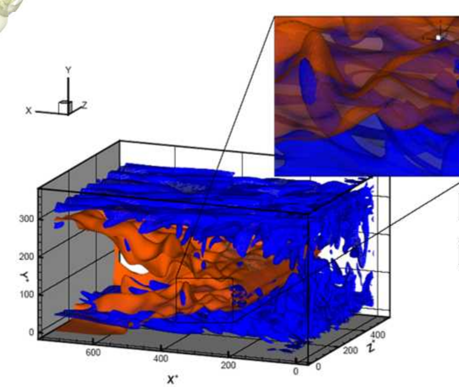

Lean premixed CH_4/air flames (2006)
200M grid points

Ethylene non-premixed sooting flames, varying Damköhler no.
350M grid points

H_2/Air Flame-wall interaction (2006)

Stratified premixed CH_4/air flames (2008)

Sandia National Laboratories

A Changing World: New Fuels

- **Fuel streams are rapidly evolving:**
  - **Heavy hydrocarbons (Oil sands, oil shale, coal)**
  - **Renewable fuels (Ethanol, biodiesel)**

- **Alternative fuels present chemical complexity, even with reduced mechanisms:**
  - **CO/H_2/Air: 11 Species, 21 Steps**
  - **CH_4/Air: 13 Species, 9 (73) Steps**
  - **C_2H_4/Air: 22 Species, 18 (167) Steps**
  - **Reduced mechanism for Biodiesel surrogates?**

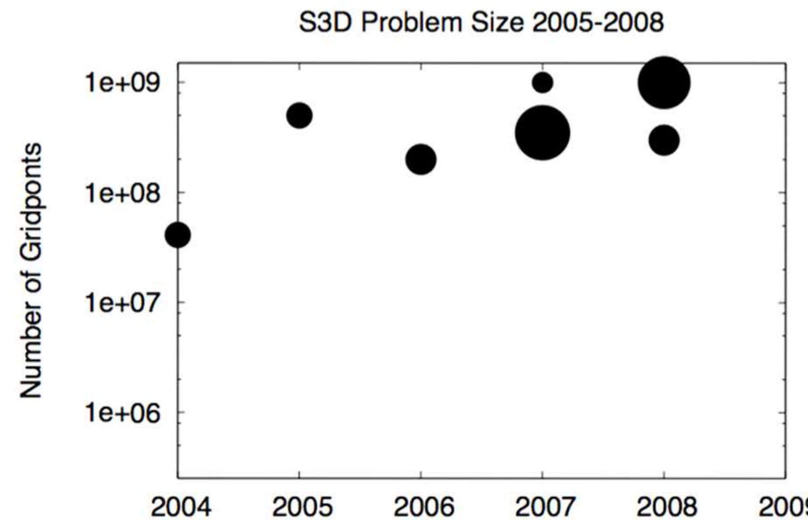

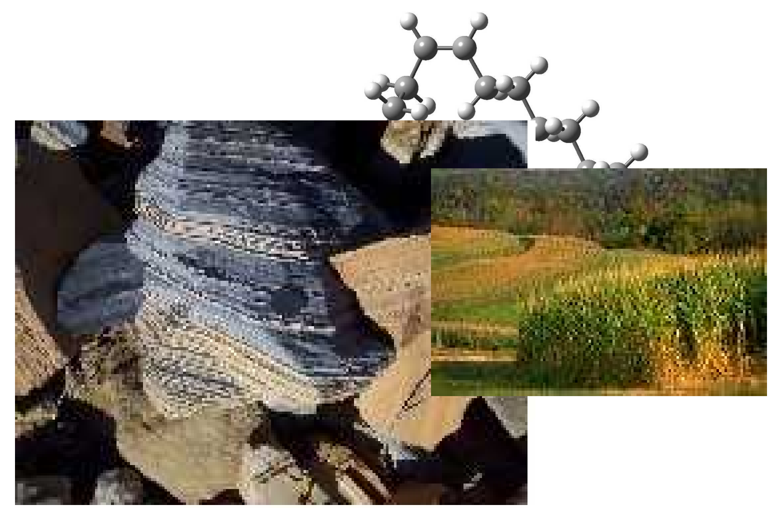

Sandia National Laboratories

Lifted Flame Stabilization

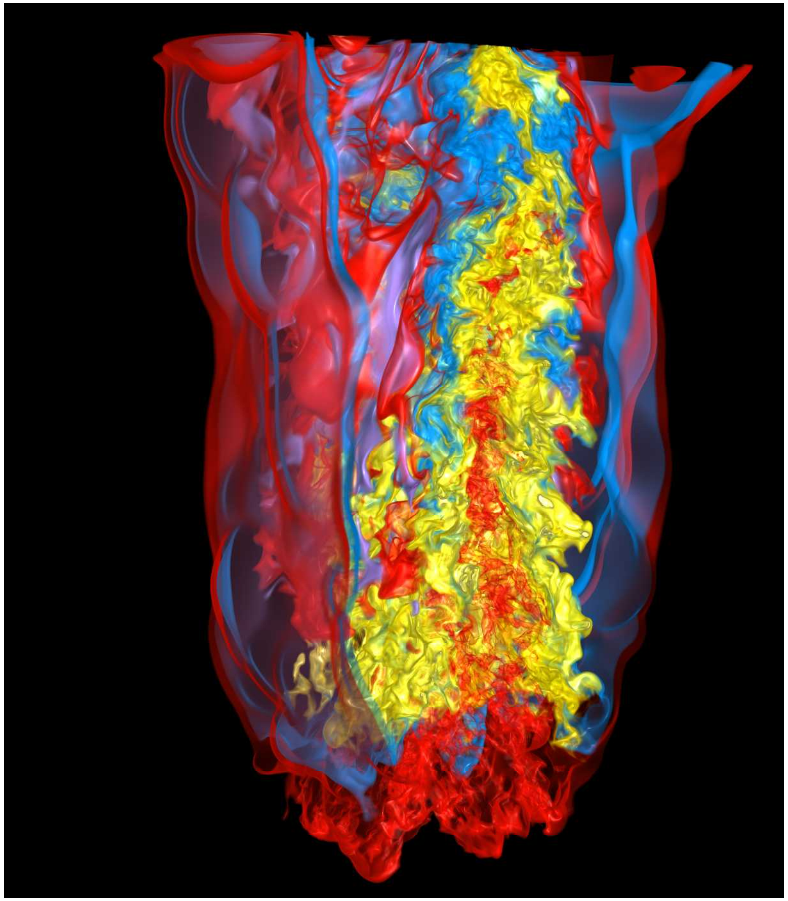

Instantaneous volume rendering of HO_2 (ignition marker, red and yellow) and OH (flame marker, blue)) by H. Yu and K.L. Ma

- **DNS reveals flame stabilizes in fuel-lean mixture where temperature is high and scalar dissipation rates are low**
- **Hydroperoxy radical (HO_2)**
  - **Chemical marker of autoignition in hydrogen-air chemistry, experimentally unobtainable**
  - **Accumulates upstream of OH and other high-temperature radicals**

2008 INCITE award enabling DNS of lifted diesel surrogate chemistry for understanding lifted diesel jet stabilization at pressure

Sandia National Laboratories

S3D Computational Cost

- Cost of computing per grid-point per time-step
  - Does not depend on number of MPI-threads
    - S3D exhibits ideal weak scaling
    - Cost of communication is minimal
    - Performance limited by single-socket compute rate
  - Depends on architecture
  - Depends on chemical mechanism (number of DOF)
  - Depends on fixed problem size per node

- Problem size varies with mechanism
  - H_2 Lifted Jet required 944 million grid points
  - C_2H_4 Lifted Jet required 1.3 billion grid points
  - Difference due to larger domain and longer residence time for ignition of hydrocarbon versus hydrogen mechanism

Sandia National Laboratories

Explicit vs Implicit Solvers for Large Mechanisms*

- Cost of implicit solvers ~ $O(K^3)$
  - Evaluation of analytic Jacobian, O(K)
  - LU decomposition of Jacobian, $O(K^3)$
- Cost of explicit solvers, O(K)
  - Chemical rates, synchronization & misc., O(K)
  - Detailed diffusion, $O(K^2$, or $K^3)$, eliminate cost with diffusive species bundling (Lu & Law 2007)
- Explicit solver is asymptotically optimal in efficiency for large mechanisms, provided stiffness can be removed

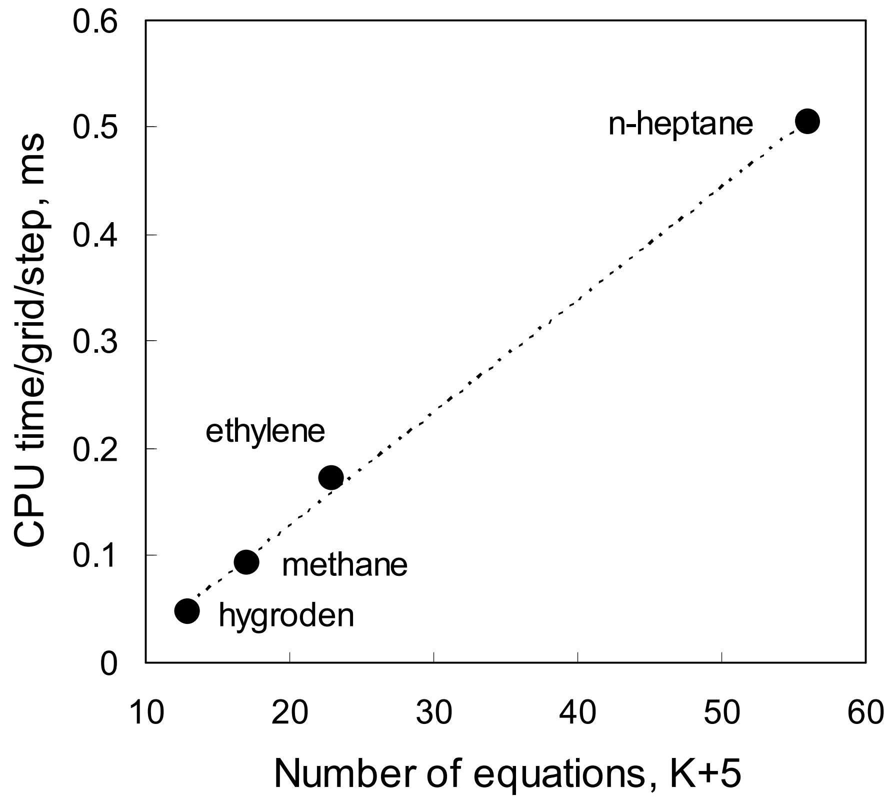

Average CPU Time for 3-D DNS (S3D) with explicit Runge-Kutta solver

*Sponsored by DOD Air Force Office of Scientific Research

Chemical Complexity vs Scale

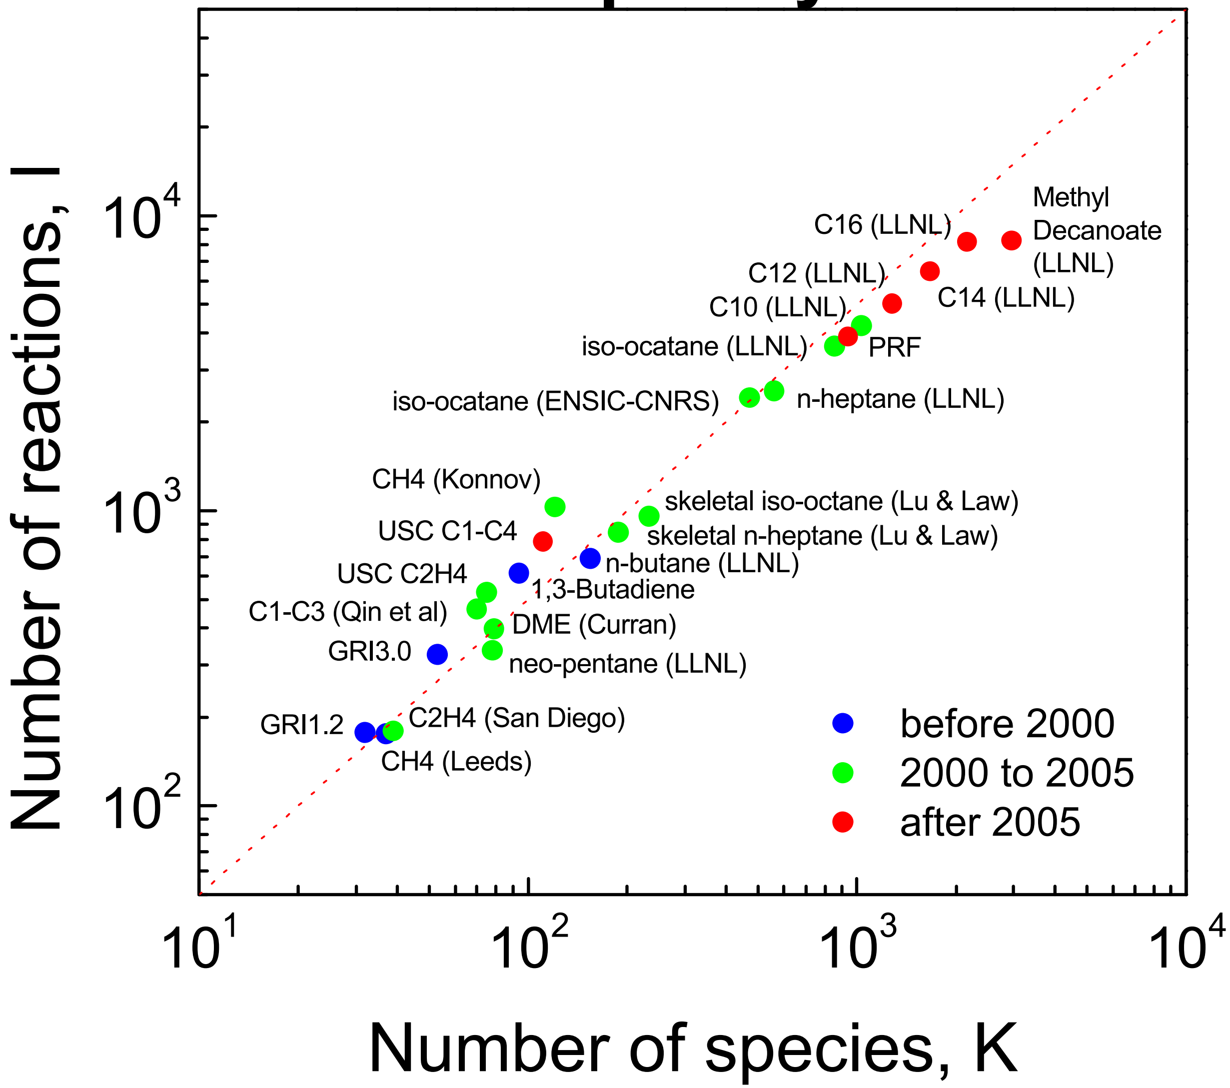

Sandia National Laboratories

Distribution of Cost is Changing

Task	H_2/Air	C_2H_4/Air
Derivatives	**15.2%**	**8.3%**
Diffusion *(Coefficients)*	**31%** *(2.4%)*	**24.6%** *(5.5%)*
Convection	**21%**	**15.3%**
Communication	**4.3%**	**4.6%**
Reaction Rates	**15.5%**	**38.2%**
I/O	**-**	**-**
Other	**13%**	**9%**

Sandia National Laboratories

2008 S3D Developments

- Driven by quad-core upgrade at ORNL
  - Chemistry routines re-written to enable compiler vectorization (John Levesque, Nathan Wichmann, & Ramanan Sankaran)
  - Using Cray provided ACML library for evaluation of exponentials with SSE vectorization
- I/O Improvements
  - Cooperative IO
  - Optimized MPI-IO library (Wei-Keng Liao)

2008 S3D Developments

- ADIOS Integration
  - Cooperative effort with Scott Klasky
- Offloading of chemistry evaluation to GPU
  - Jeffrey Vetter, Jeremy Meredith, Kyle Spafford
  - Exploratory use of heterogeneous architectures
  - Double precision / single precision
- Subversion repository
  - https://outrage.ca.sandia.gov/svn/s3d
  - Ethylene test case will be checked in next week

Rewrite of getrates() for C_2H_4

- Rewrite to use f90 array syntax:
  - getrates(T,P,Y(:),W) -> getrates(T(:), P(:), Y(:,:),W(:))
  - Compiler based SSE optimization
  - Cross platform benefit
  - Auto-generation utilities updated to output new syntax for skeletal mechanisms (e.g. H_2/Air)
- Rewrite to use ACML exponential evaluation
  - RF(:,i) = exp(RF(:,i) -> vrda_exp(s,RF(1,i),RF(1,i))
  - Benefit to all platforms where ACML available
  - Cray provided ACML library available on XT4

Distribution of Cost

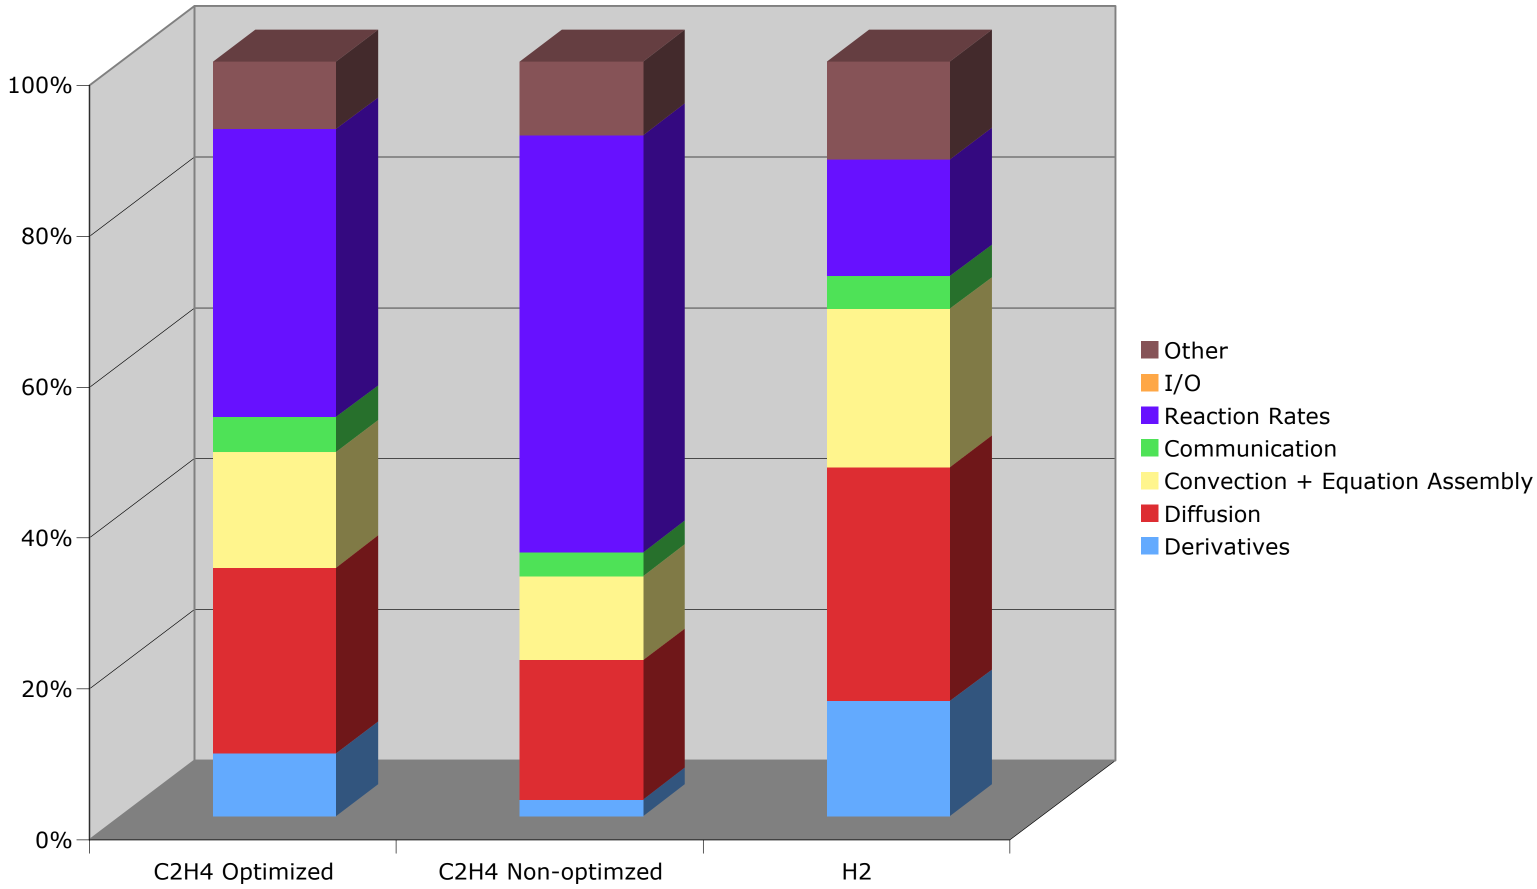

Sandia National Laboratories

Summary

- Weak scaling nature of S3D has allowed on-par increase in problem size with compute power
- To continue to fulfill S3D's critical role, density of point-wise compute burden will increase
- Chemistry routine is still low-risk target for optimization / offload to coprocessor
- We need guidance on restructuring the code for efficiency gains – asynchronous communication and deep memory hierarchies
- As memory requirement increases with transported species, restructure big arrays to maximize data reuse

Quad core performance

- Weak scaling test with C_2H_4 mechanism.
  - 30^3 grid points per MPI thread on XT4
  - 100 time steps
- 2.6 GHz dual-core Opteron and 2 Flops per cycle
- 2.1 GHz quad-core Opteron and 4 Flops per cycle

Fixed Problem Size	MPI Mode	Dual	Quad	Quad (Improved)
		μseconds per gridpoint per timestep	μseconds per gridpoint per timestep	μseconds per gridpoint per timestep
30x30x30	aprun -n 1 -N 1	150	154	123
60x30x30	aprun -n 2 -N 2	172	159	129
60x60x30	aprun -n 4 -N 4	-	186	156

On a per-core basis the quad-core is performing faster than the dual-core system

Sandia National Laboratories

Effect of vectorization

Times from quad-core	Before	After
	seconds	
Rates time	277	228
Total time	415	367

Counters from dual-core	Before	After
	$\times 10^9$ Operations	
Add	182	187
Multiply	204	210
Add+Multiply	386	397
Load/Store	179	202
SSE	91	212