SAND2014-15368PE

Exceptional service in the national interest

Sandia National Laboratories

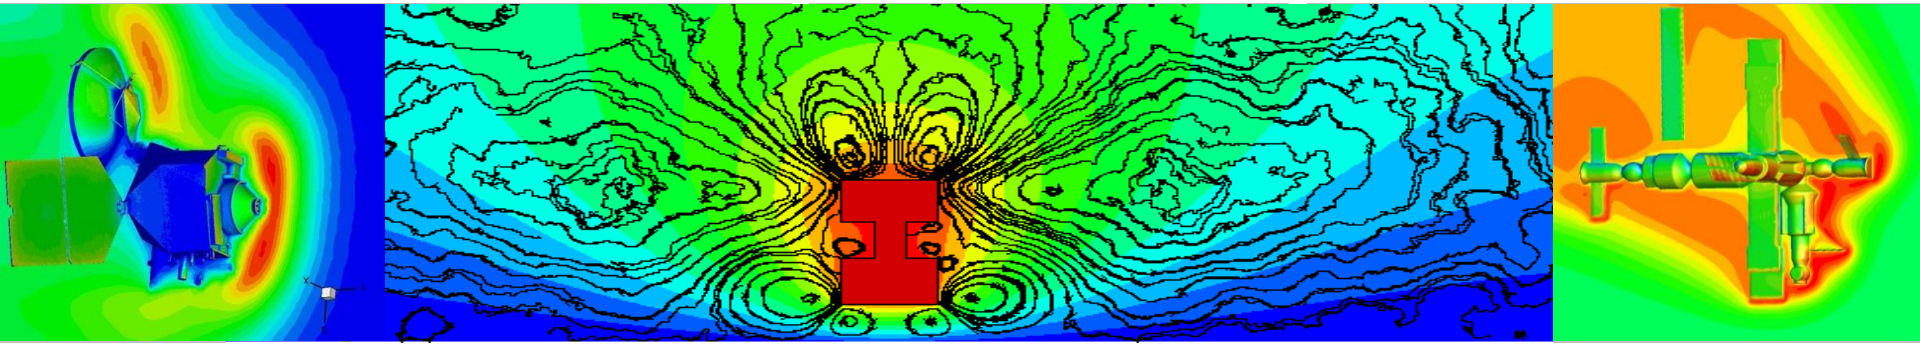

DSMC: The Quest for Speed

Michael A. Gallis

Engineering Sciences Center
Sandia National Laboratories
Albuquerque, New Mexico, USA

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Collaborators

Sandia National Laboratories


John Torczynski


Dan Rader


Tim Koehler


Steve Plimpton

Boltzmann Equation and the Direct Simulation Monte Carlo Method

Sandia National Laboratories

Ludwig Boltzmann

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{x}} + \frac{\mathbf{F}}{m} \cdot \frac{\partial f}{\partial \mathbf{v}} = \int_{-\infty}^{\infty} \int_{0}^{4p} \left(f^* f_1^* - f f_1 \right) \left| \mathbf{v} - \mathbf{v}_1 \right| s \, dW d\mathbf{v}_1$$

molecular motion and force-induced acceleration

pairwise molecular collisions (molecular chaos)

James Clerk Maxwell

The objective of DSMC is to simulate complicated gas flows using only collision mechanics of simulated molecules

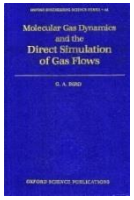

Graeme Bird (1963, 1994)

molecules move **molecules collide**

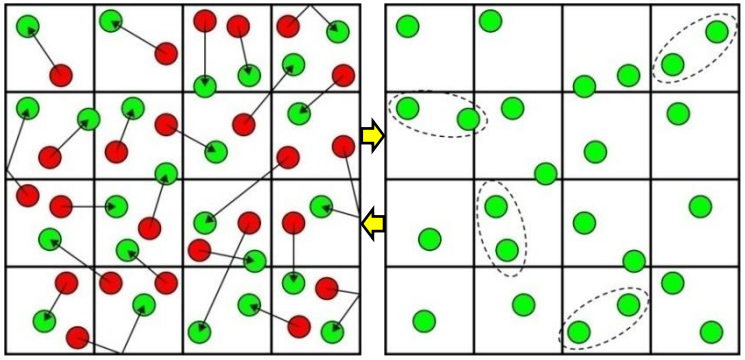

DSMC vs. Boltzmann Equation

Sandia National Laboratories

- Instead of solving Newton's laws of motion (Molecular Dynamics), DSMC replaces explicit intermolecular forces with stochastic collisions
- It has been shown that DSMC is equivalent to solving the Boltzmann equation (Nambu 1980, Babovsky 1989, Wagner 1992)
- DSMC has been shown to reproduce exact known solutions (Chapman-Enskog, Moment Hierarchy) of the Boltzmann equation (Gallis et al. 2004, 2006) for non-equilibrium flows
- In fact, DSMC is superior to solving the Boltzmann equation
  - DSMC can model complicated processes (e.g., polyatomic molecules, chemically reacting flows, ionized flows) for which Boltzmann-type transport equations are not even known (Struchtrup 2005)
  - DSMC includes fluctuations, which have been shown to be physically realistic (Garcia 1990) but which are absent from the Boltzmann equation
- DSMC has tremendous potential
- 1000x speedup is desired to address problems of interest

Evolution of DSMC Procedures

Sandia National Laboratories

- Since the original DSMC algorithm (1963), there have been no major changes to it
  - Advanced molecular interactions, physical phenomena have been added
- The DSMC algorithm has been criticized (unfairly) as being computationally inefficient
- Numerous alternative algorithms or procedures have been proposed that try to address this concern
- Alternative methods come with undeniable virtues but accompanied vices.
- There is no standardized test suite to allow comparative evaluations of new methods, which has led to a proliferation of *new methods*
- More than 50 years since its introduction, DSMC is still the predominant algorithm

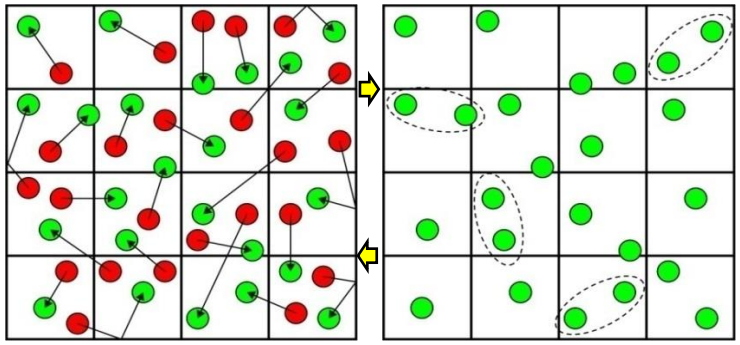

Alternative Procedures

Sandia National Laboratories

- Plume flows
  - Hybrid schemes (NS-DSMC)
- Non-reacting flows
  - Discrete velocity schemes, moment methods, higher-order CFD
- Low signal/noise flows (MEMS)
  - PDE-like, numerical solutions of the Boltzmann equation
  - Noise reduction schemes
  - Simplified molecular interactions (BGK/ES-BGK)
- High-density flows
  - Near-neighbor collisions
- Free-molecular flows
  - Analytical solutions

How can we evaluate new techniques & procedures?

Sandia National Laboratories

- There is a very large number of alternative procedures with an ever-growing number of variants
- Although there are particular applications where some variants may be preferable, it is unclear how to evaluate
  - Generality
  - Accuracy
  - Efficiency

- Accuracy → Comparison with analytical solutions
- Efficiency → Comparison with DSMC convergence rate

Chapman-Enskog (CE) Theory

Sydney Chapman

David Enskog

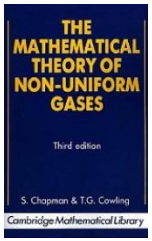

$$f = f^{(0)}(1 + \mathrm{F}^{(1)} + \mathrm{Y}^{(1)}) \qquad f^{(0)} = (n/\pi^{3/2} c_m^3)\exp[-\tilde{c}^2]$$

$$c_m = \sqrt{2k_B T/m} \qquad \tilde{\mathbf{c}} = \mathbf{c}/c_m \quad \mathbf{c} = \mathbf{v} - \mathbf{u}$$

$$\mathrm{F}^{(1)} = -\left(8/5\right)\tilde{A}[\tilde{c}]\tilde{\mathbf{c}}\cdot\tilde{\mathbf{q}} \qquad \mathrm{Y}^{(1)} = -2\tilde{B}[\tilde{c}](\tilde{\mathbf{c}}\circ\tilde{\mathbf{c}} : \tilde{I})$$

$$\kappa = -(5/4)k_B c_m^2 a_1 \qquad \eta = (1/2)mc_m^2 b_1$$

$$\tilde{A}[\tilde{c}] = \sum_{k=1}^{\infty} (a_k/a_1) S_{3/2}^{(k)}[\tilde{c}^2] \qquad \tilde{B}[\tilde{c}] = \sum_{k=1}^{\infty} (b_k/b_1) S_{5/2}^{(k-1)}[\tilde{c}^2]$$

$$C_p = (5/2)(k_B/m) \qquad \mathrm{Pr} = (2/3)(\eta_\infty/\eta_1)(K_1/K_\infty)$$

- Chapman and Enskog analyzed Boltzmann collision term
  - Perturbation expansion using Sonine polynomials
  - Near equilibrium, appropriate in continuum limit
- Determined velocity distribution and transport properties
  - Thermal conductivity K, viscosity μ, mass self-diffusivity D
  - Prandtl number Pr from "infinite-to-first" ratios K_∞/K_1, μ_∞/μ_1
  - Distribution "shape": Sonine polynomial coeffs. a_k/a_1, b_k/b_1
  - Values for all Inverse-Power-Law (IPL) interactions
    - Maxwell and hard-sphere are special cases

Extracting CE Parameters from DSMC

Sandia National Laboratories

$$q = K_{eff}\left(\frac{\partial T}{\partial x}\right) \qquad \frac{a_k}{a_1} = \sum_{i=1}^{k}\left(\frac{(-1)^{i-1}k!(5/2)!}{(k-i)!i!(i+(3/2))!}\right)\left(\frac{\left\langle \tilde{c}^{2i}\tilde{c}_x\right\rangle}{\left\langle \tilde{c}^{2}\tilde{c}_x\right\rangle}\right) \qquad \tilde{\mathbf{c}} = \frac{\mathbf{v}-\mathbf{V}}{c_m}$$

$$\tau = \mu_{eff}\left(\frac{\partial V}{\partial x}\right) \qquad \frac{b_k}{b_1} = \sum_{i=1}^{k}\left(\frac{(-1)^{i-1}(k-1)!(5/2)!}{(k-i)!(i-1)!(i+(3/2))!}\right)\left(\frac{\left\langle \tilde{c}^{2(i-1)}\tilde{c}_x\tilde{c}_y\right\rangle}{\left\langle \tilde{c}_x\tilde{c}_y\right\rangle}\right) \qquad c_m = \sqrt{\frac{2k_BT}{m}}$$

DSMC moments of velocity distribution function

- Temperature T, velocity V
- Heat flux q, shear stress τ
- Higher-order moments

DSMC values for VSS molecules (variable-soft-sphere)

- Thermal conductivity and viscosity: K_{eff} and μ_{eff}
- Sonine-polynomial coefficients: a_k/a_1 and b_k/b_1
- Applicable for arbitrary Kn_L, Kn_q, Kn_τ

Fourier and Couette Flow

Joseph Fourier

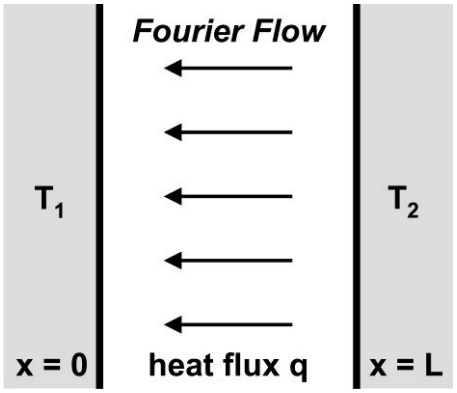

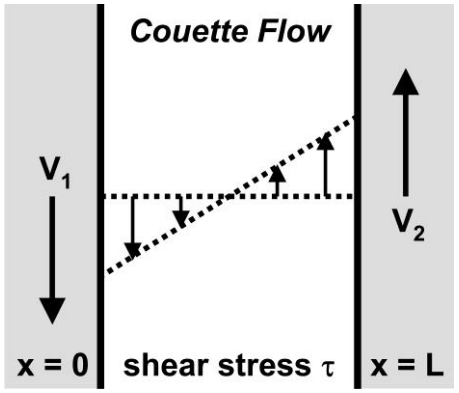

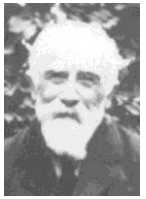

Maurice Couette

$$q = -K\frac{\partial T}{\partial x}$$

$$\tau = \mu\frac{\partial v}{\partial x}$$

Investigate transport in gas between parallel plates

- Fourier flow: heat conduction in stationary gas
- Couette flow: momentum transport in isothermal shear flow

Apply DSMC to Fourier flow and Couette flow

- Heat flux, shear stress: one-dimensional, steady

Compare DSMC to analytical "normal solutions"

- Normal: outside Knudsen layers
- Solutions: Chapman-Enskog (CE), Moment-Hierarchy (MH)

Verify DSMC accuracy at arbitrary heat flux, shear stress

- Thermal conductivity, viscosity; velocity distribution

Temperature and Velocity Profiles

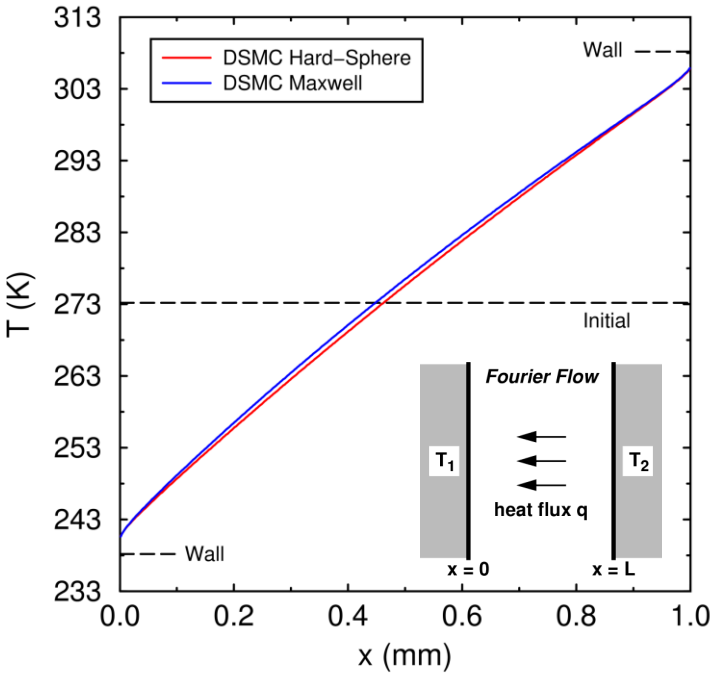

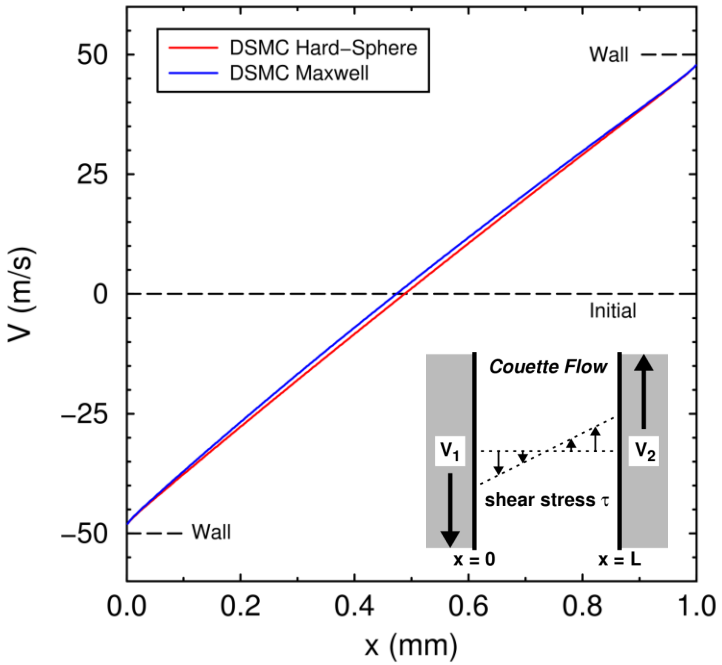

Low heat flux and shear stress: $Kn_q = 0.006$, $Kn_\tau = 0.003$

- Argon-like: initial T = 273.15 K, p = 266.644 Pa, λ = 24 μm
- Walls: L = 1 mm = 42λ, ΔT = 70 K, ΔV = 100 m/s
- N_c = 120, Δt = 7 ns, Δx = 2.5 μm, ~10^9 samples/cell, 32 runs

Small velocity slips, temperature jumps, Knudsen layers

DSMC Reproduces Infinite-Approximation Chapman-Enskog Transport Coefficients

Sandia National Laboratories

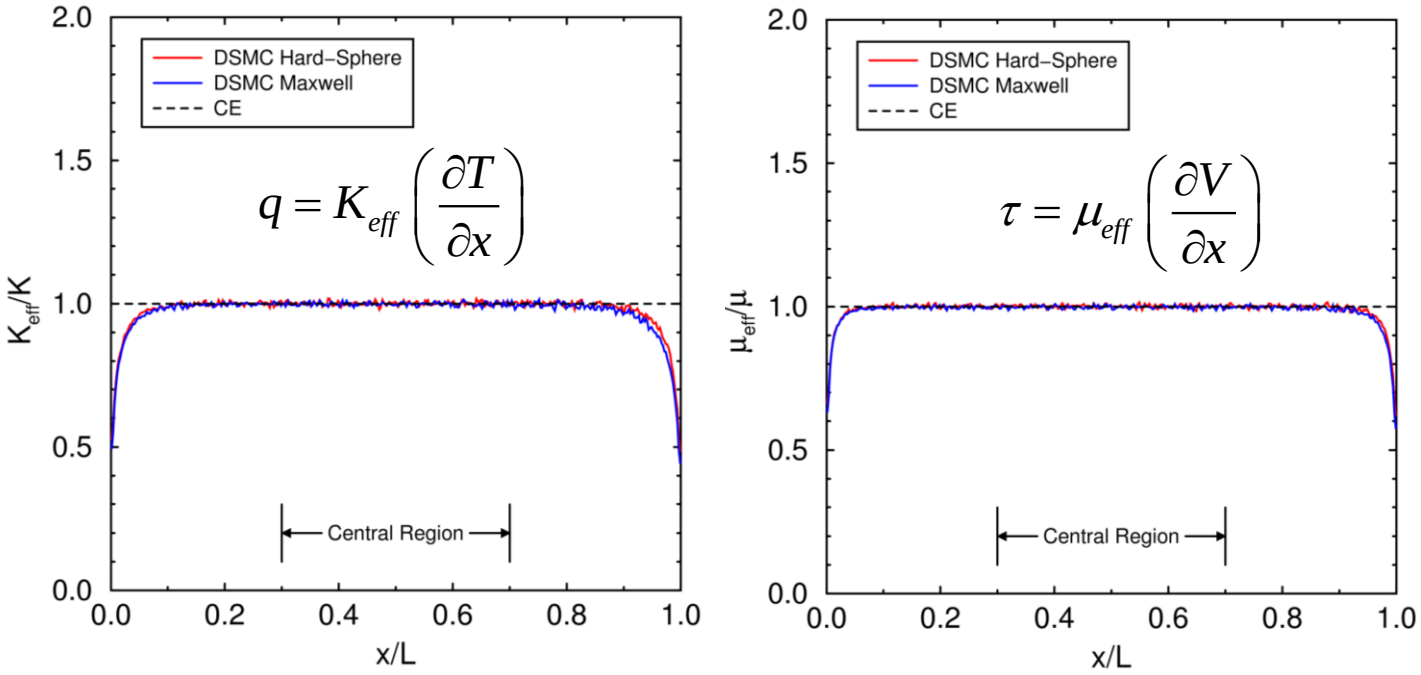

Thermal conductivity (left) and viscosity (right) away from walls

- Maxwell and hard-sphere results bound most gases
- Agreement with Chapman-Enskog theory verifies DSMC

DSMC Reproduces Infinite-Approximation Chapman-Enskog Velocity Distribution

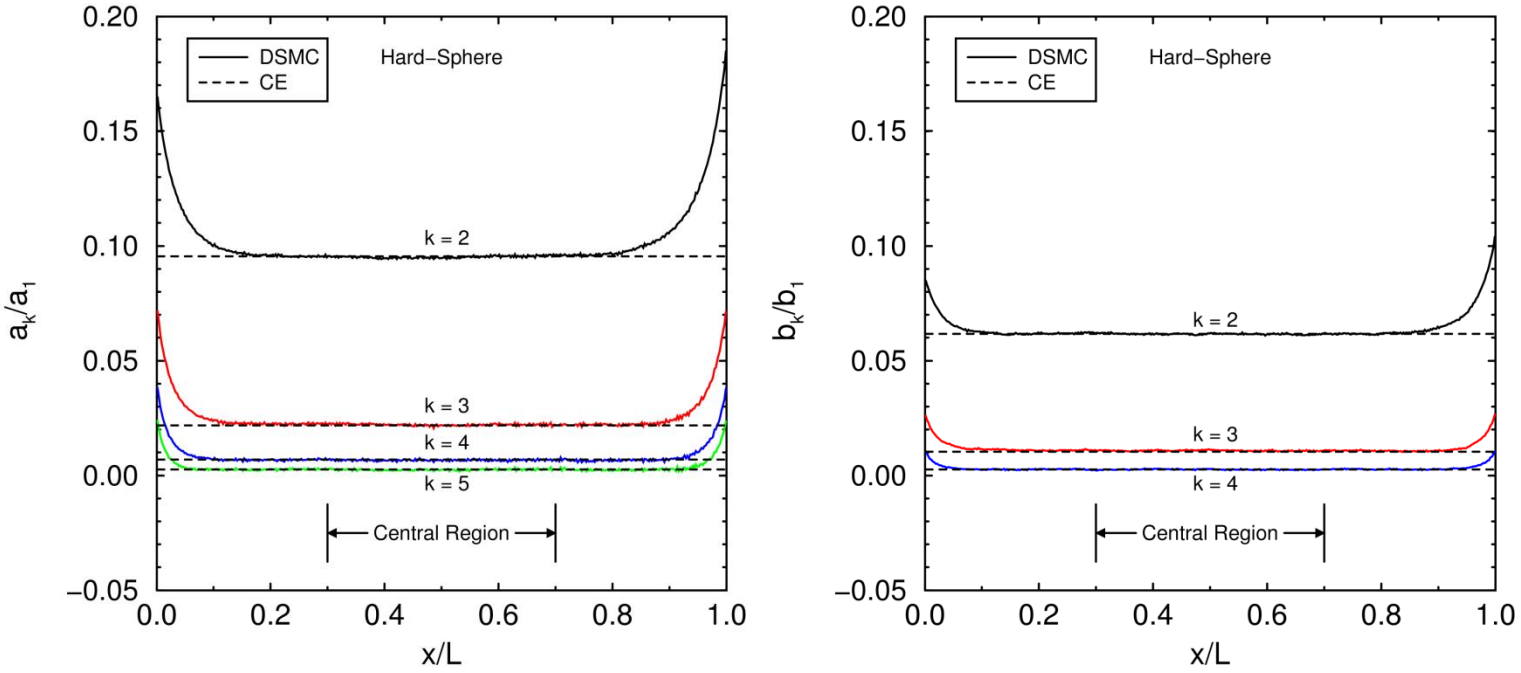

Sonine polynomial coefficients for temperature (left) & velocity (right) gradients

- **Hard-sphere values are shown, other interactions have similar agreement**
- **Higher-order (k > 5) coefficients (not shown) also have similar agreement**

Gallis M. A., Torczynski J. R., Rader D. J., "Molecular Gas Dynamics Observations of Chapman-Enskog Behavior and Departures Therefrom in Nonequilibrium Gases", *Physical Review E*, 69, 042201, 2004.

Maxwell Sonine-Coefficient Profiles

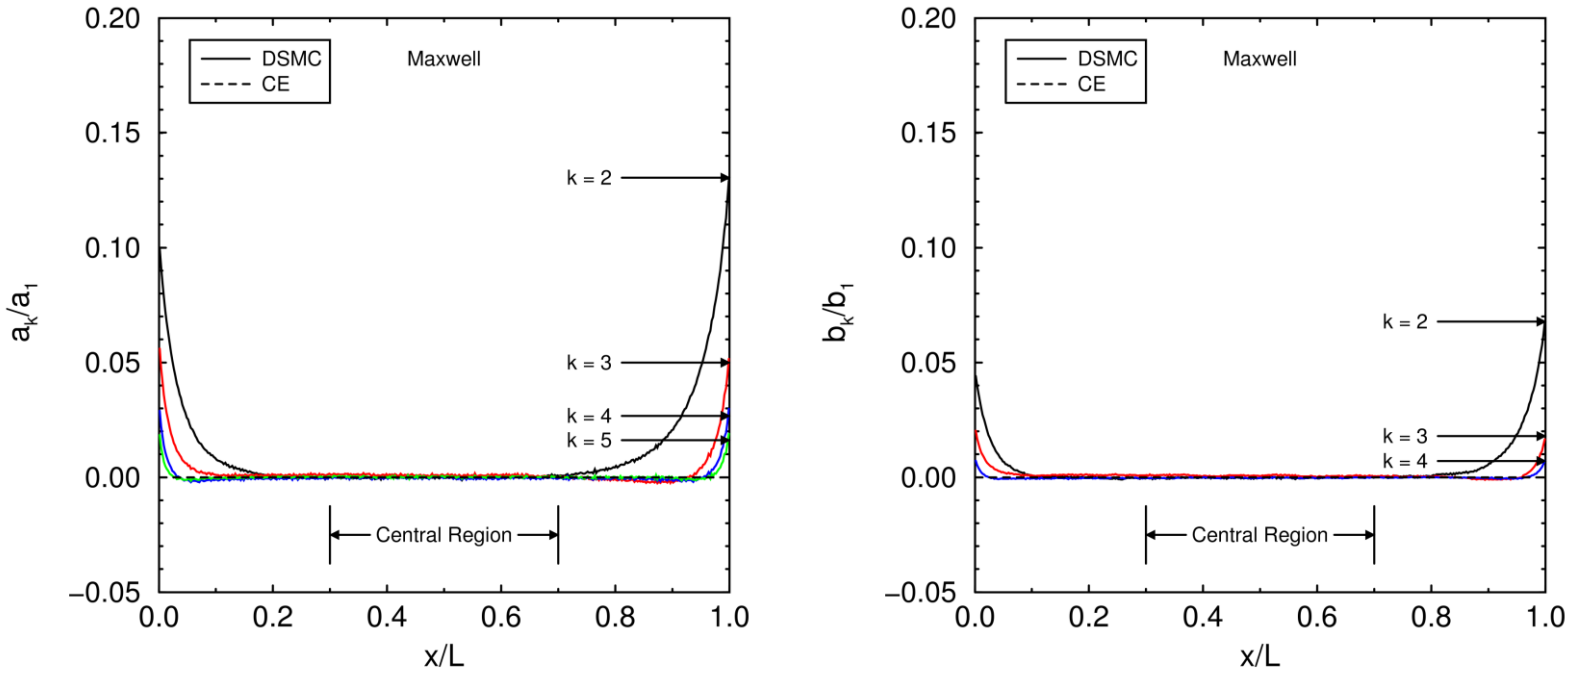

DSMC and CE Maxwell coefficients a_k/a_1 and b_k/b_1

- Low heat flux, low shear stress: $Kn_q = 0.006$, $Kn_\tau = 0.003$
- Good agreement in central region: normal solution
- Knudsen layers easily observed: ~10% of domain

Moment-Hierarchy Method

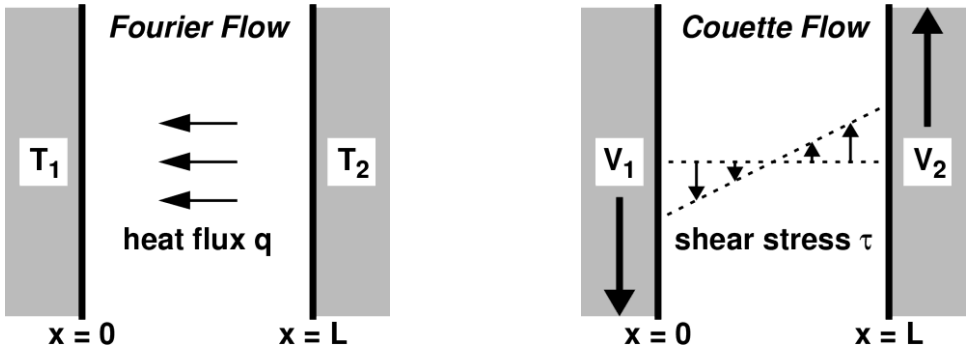

Moment-Hierarchy (MH) normal solution

- MH solution extends CE solution to finite Kn_q and Kn_t
- Maxwell molecules: collision term quadratic in moments

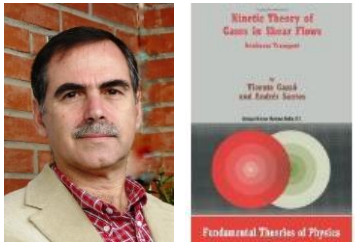

Andres Santos

Compare DSMC to MH for Maxwell molecules

- Dependence of Sonine coefficients on Kn_q known

Apply DSMC for Maxwell molecules

Gallis M. A., Torczynski J. R., Rader D. J., Tij M., Santos A., "Normal Solutions of the Boltzmann Equation for Highly Nonequilibrium Fourier and Couette Flow", *Phys. Fluids, 18, 017104,* 2006.

Maxwell Normalized Sonine Coefficients

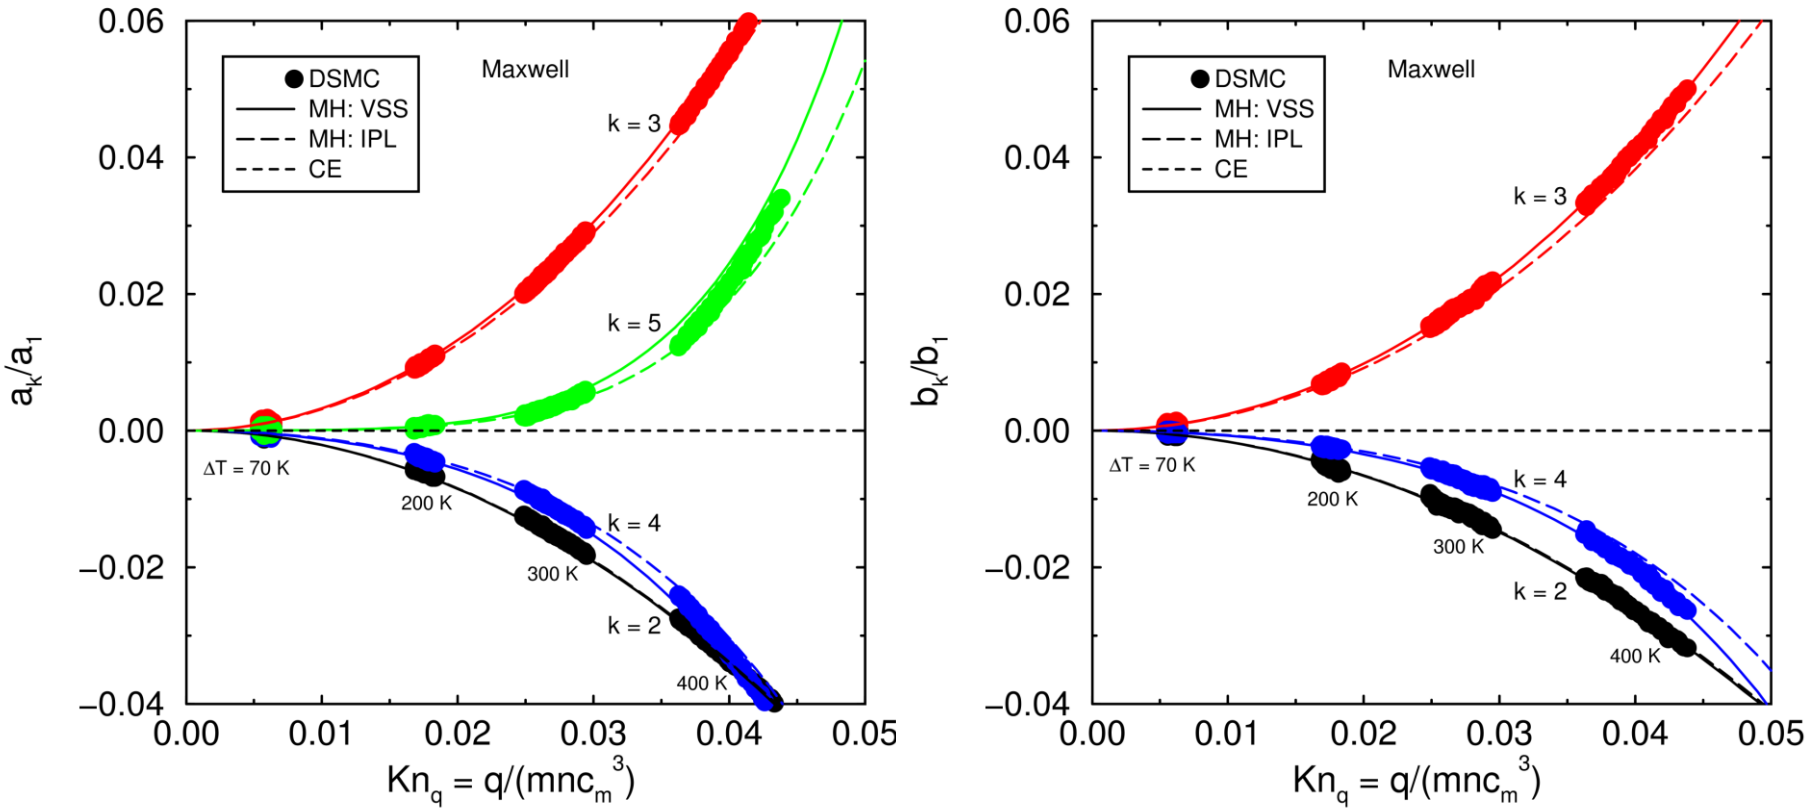

DSMC and MH Maxwell normal solutions for a_k/a_1 and b_k/b_1

- Four DSMC simulations: ΔT = 70, 200, 300, 400 K
- MH: VSS-Maxwell (solid) and IPL-Maxwell (dashed) differ
- DSMC and MH VSS-Maxwell normal solutions agree

ES-BGK Kinetic Model

Lowell Holway

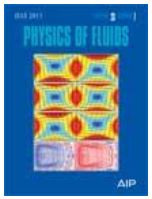

Carlo Cercignani

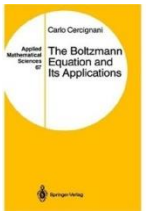

Ellipsoidal-Statistical Bhatnagar-Gross-Krook
Holway (1963, 1965, 1966); Cercignani (1967, 1988)

$$\left[\frac{df}{dt}\right]_{\mathrm{ES}} = \nu_{\mathrm{ES}}\left(f_{\mathrm{ES}} - f\right) \qquad f_{\mathrm{ES}} = \frac{n}{\pi^{3/2} c_m^3 \det\left[\mathbf{S}\right]} \exp\left[-\tilde{\mathbf{C}}^{\mathrm{T}}\tilde{\mathbf{C}}\right]$$

$$\tilde{\mathbf{c}} = \mathbf{S}\tilde{\mathbf{C}} \qquad \mathbf{P} = \int \mathbf{c}\mathbf{c}^{\mathrm{T}} f \, d\mathbf{v} \qquad \mathbf{SS} = \mathbf{I} - \left(\frac{1}{\mathrm{Pr}} - 1\right)\left(\frac{3\mathbf{P}}{\mathrm{trace}\left[\mathbf{P}\right]} - \mathbf{I}\right)$$

$$K_{\mathrm{ES}} = nk_B T C_p / \nu_{\mathrm{ES}} \qquad \mu_{\mathrm{ES}} = nk_B T \,\mathrm{Pr} / \nu_{\mathrm{ES}}$$

$$C_p = (5/2)(k_B/m) \qquad \mathrm{Pr}_{\mathrm{ES}} = \mathrm{Pr}$$

ES-BGK collision term is generalization of BGK collision term

- Replaces Maxwellian with Ellipsoidal-Statistical distribution
- Prandtl number is specified independently (say, ~2/3)
- Drives distribution to equilibrium (Andries et al., 2000)

ES-BGK can presumably simulate a broad range of flows

- Can match thermal conductivity and viscosity simultaneously

Gallis M. A., Torczynski J. R., "Investigation of the ellipsoidal-statistical Bhatnagar-Gross-Krook kinetic model applied to gas-phase transport of heat and tangential momentum between parallel walls", Phys. Fluids, 23(3), 030601, 2011.

Distribution Shape: Maxwell

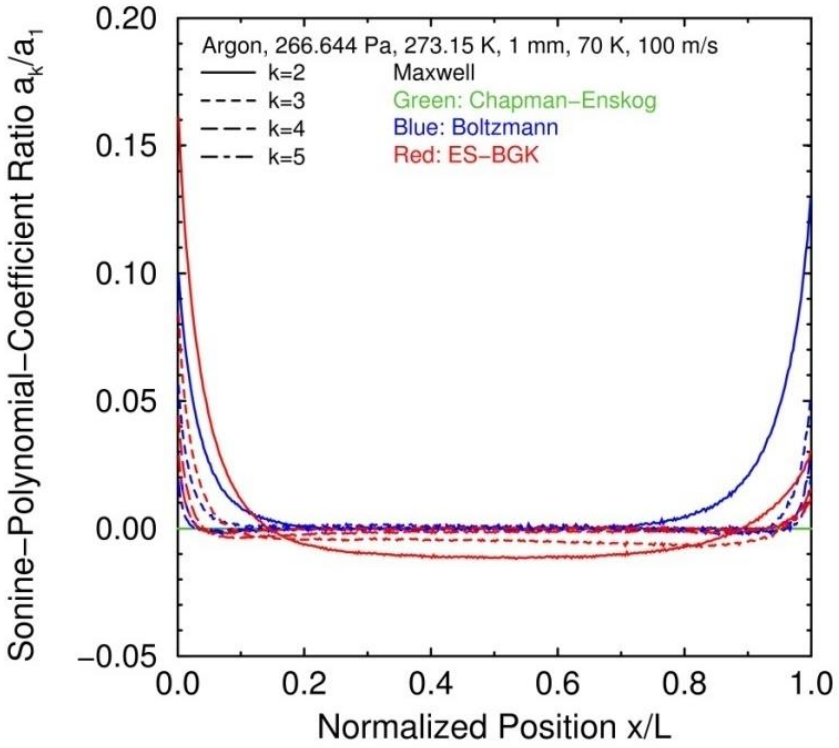

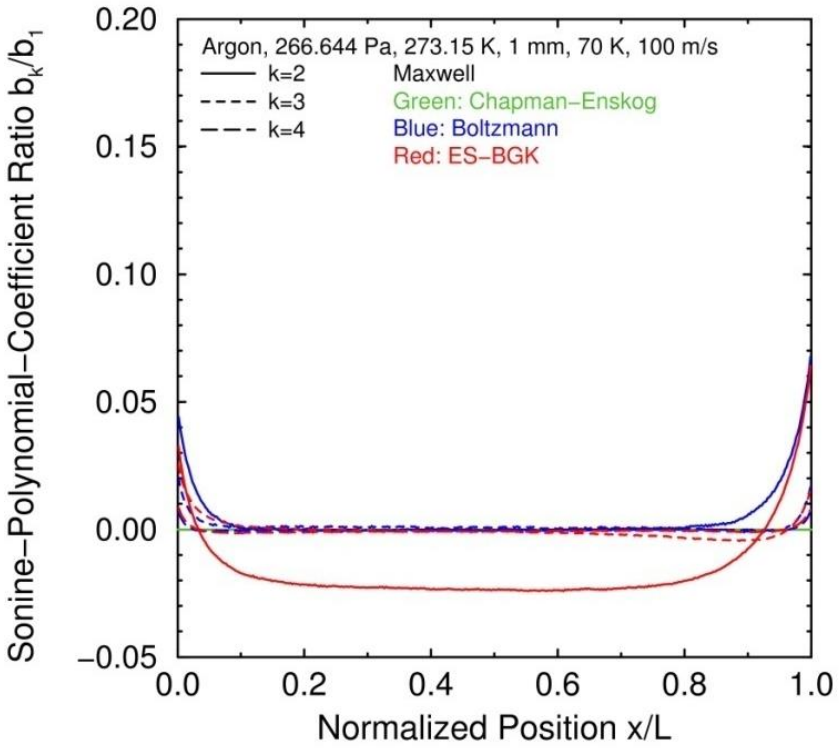

Sonine polynomial coefficients for Maxwell interaction

- Chapman-Enskog (continuum) values for k ≥ 2 are all zero
- Boltzmann closely matches CE values away from walls
  - Systematic differences in Knudsen layers (transition regime)
- ES-BGK differs strongly from CE values away from walls
  - Discrepancy is largest for k = 2

ES-BGK distribution shape is not accurate in continuum and transition regimes

Distribution Shape: Hard-Sphere

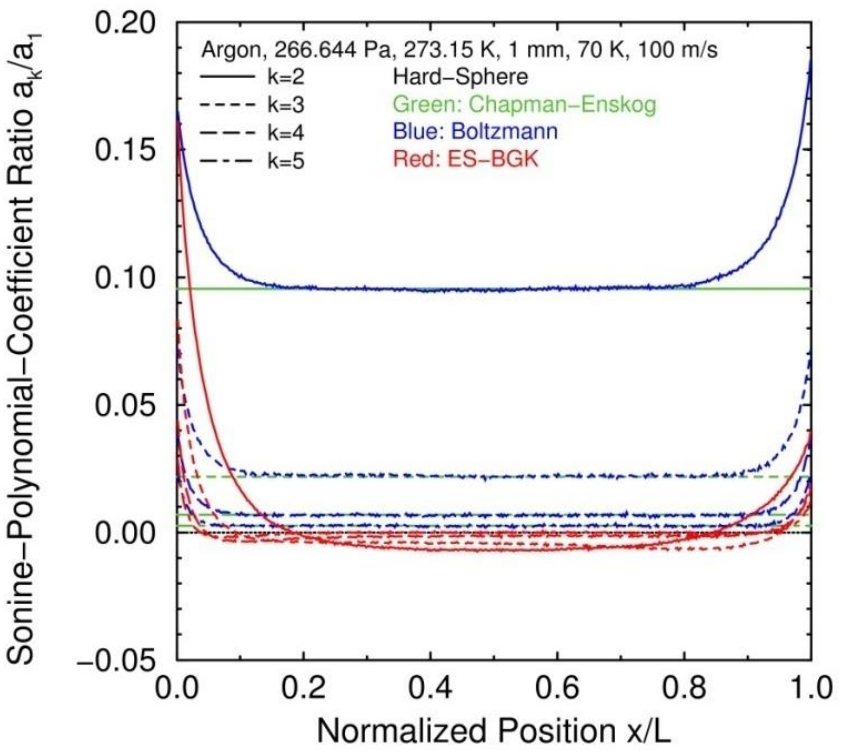

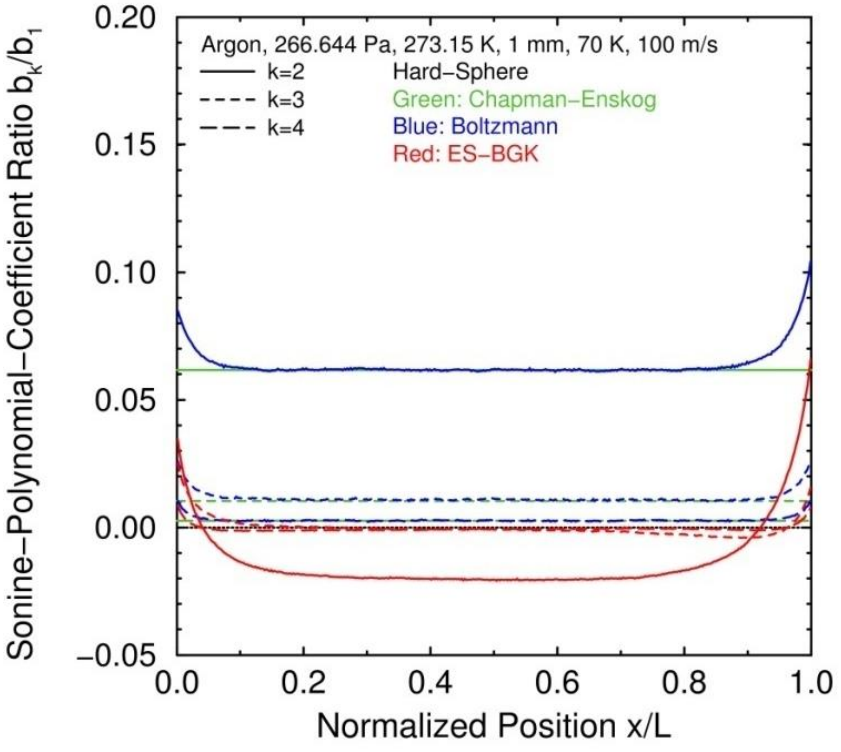

Sonine polynomial coefficients for hard-sphere interaction

- Chapman-Enskog (continuum) values for k ≥ 2 are nonzero
  - Systematic differences in Knudsen layers (transition regime)
- ES-BGK differs strongly from CE values away from walls
  - ES-BGK hard-sphere & Maxwell values hardly differ but should

Variation of ES-BGK distribution shape with molecular interaction is not accurate

Distribution Shape: Maxwell

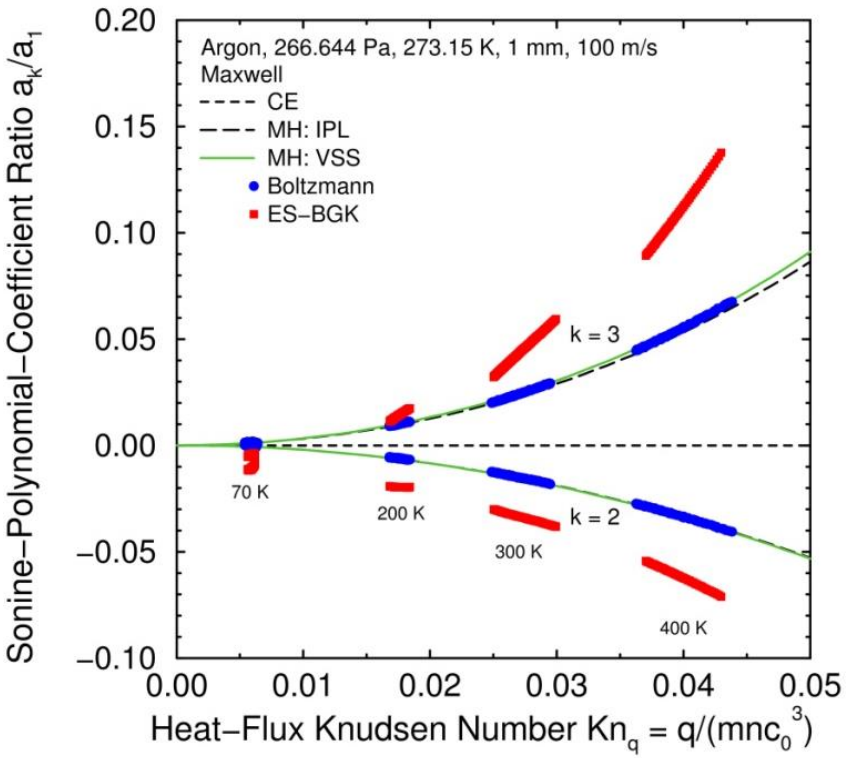

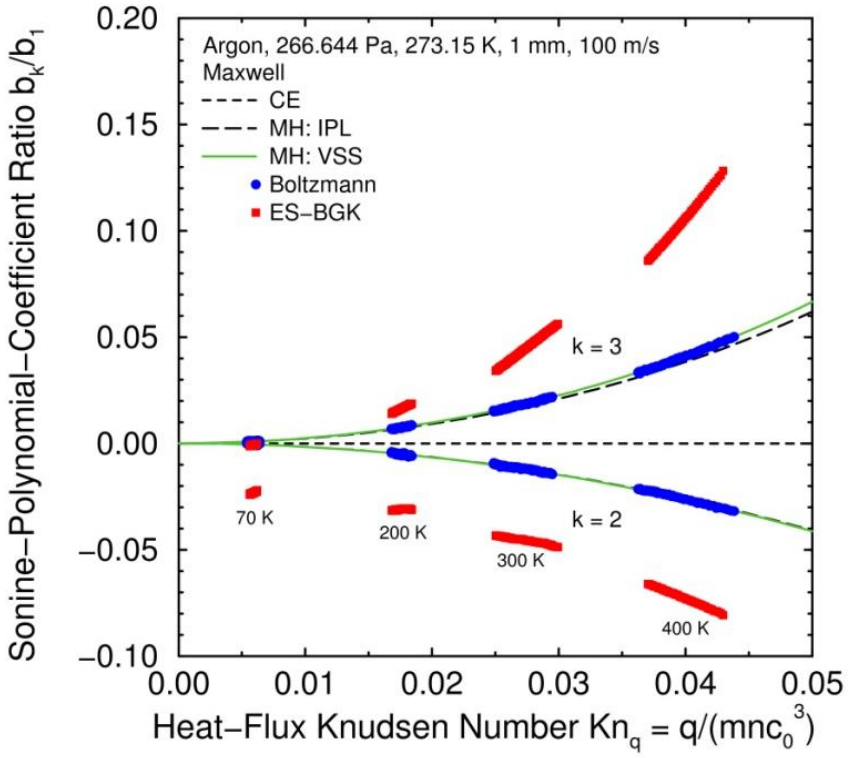

- ES-BGK shape is inaccurate at large flux or stress
- Chapman-Enskog (CE) values for k ≥ 2 are all zero
  - Appropriate only in limit of zero heat flux & shear stress
- Moment-Hierarchy (MH) values for k ≥ 2 are nonzero
  - Appropriate for arbitrary nonzero heat flux & shear stress

Fickian Mass Diffusion

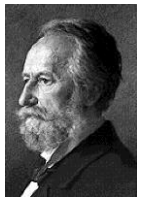

Adolf Fick

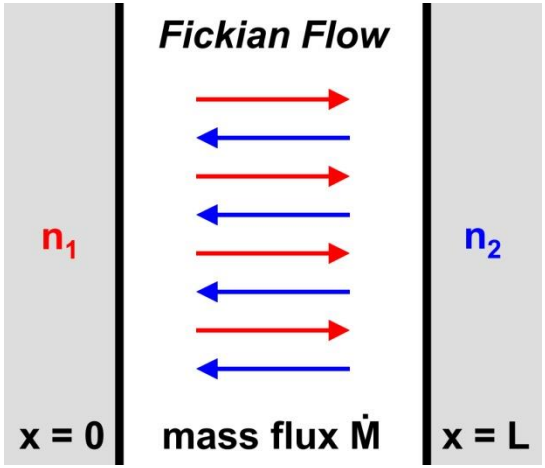

$$D = \frac{\mu}{\rho}\frac{D_\infty / D_1}{\mu_\infty / \mu_1}\frac{3(3\nu-5)}{5(\nu-1)}\frac{A_2[\nu]}{A_1[\nu]}$$

$$\nu = 1 + \frac{4}{2\omega - 1} \qquad n_1 u_1 + n_2 u_2 = 0$$

$$n_1 u_1 = -D\frac{\partial n_1}{\partial x} \qquad n_2 u_2 = -D\frac{\partial n_2}{\partial x}$$

Investigate transport in gas between parallel plates

- Fickian flow: mass diffusion in stationary isothermal gas
  - Identical molecules are passively tagged either “1” or “2”
  - Reflection sets tag to “1” at left wall and “2” at right wall

Compare Boltzmann, ES-BGK results

- Number density, velocity, and mass self-diffusivity profiles
- Maxwell and hard-sphere molecular interactions
- Continuum cases and Chapman-Enskog theory

Mass Self-Diffusivity

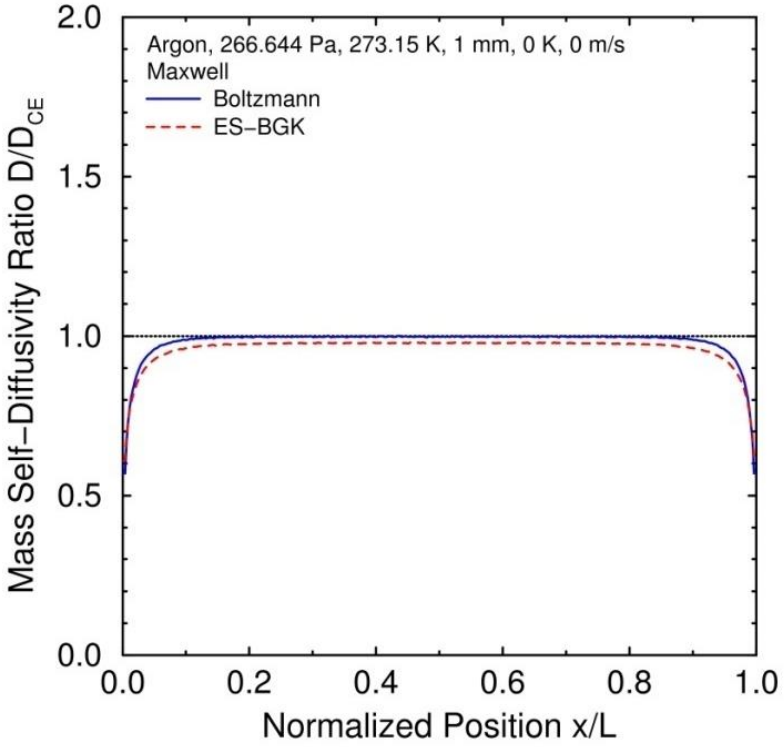

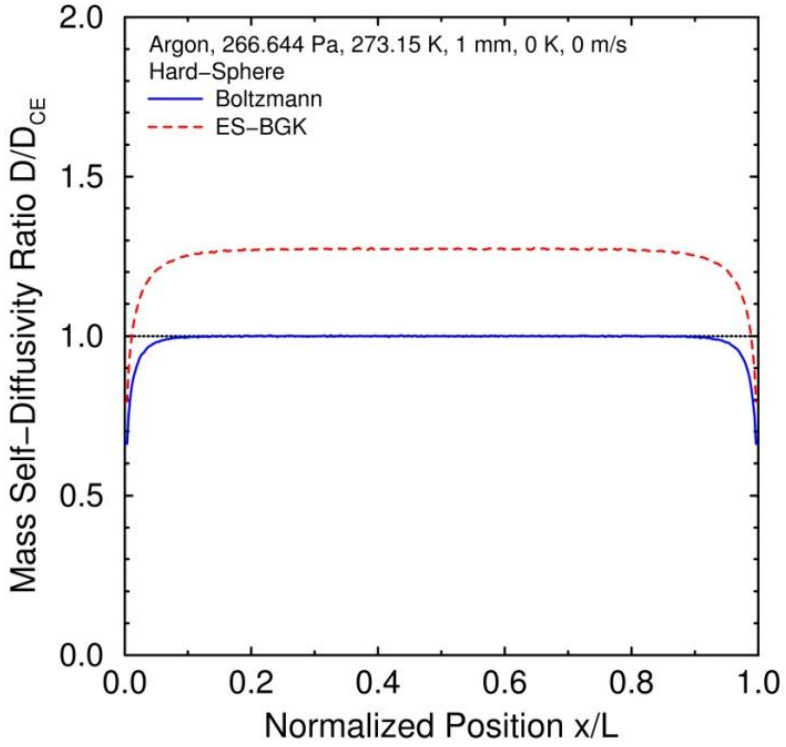

ES-BGK is not accurate for mass diffusion

- Normalize simulation values by Chapman-Enskog values
- Boltzmann agrees well with CE values away from walls
- ES-BGK differs from CE values even away from walls
  - Slightly low for Maxwell, quite high for hard-sphere

Evaluating Computational Efficiency

Sandia National Laboratories

- Computational efficiency is more than CPU time.
  - CPU time to achieve a particular level of accuracy
  - Convergence rate
  - Parallel efficiency (hardly ever addressed)

- Comparing low-order moments (density, temperature) is usually an insensitive measure of comparison.

DSMC Numerical Error

Sandia National Laboratories

Traditional DSMC rule-of-thumb guidelines:

- Take enough samples to drive statistical error down to "acceptable" level
- Keep time step smaller than ~1/4 mean collision time
- Keep cell size smaller than ~1/3 mean free path
- Use a minimum of ~20 particles per cell

These guidelines give 2% error, which is similar to the uncertainty in measured transport properties for most gases

- DSMC is subject to the same constraints as other numerical methods.
- DSMC is correct to the limit of vanishing discretization.

DSMC Numerical Error

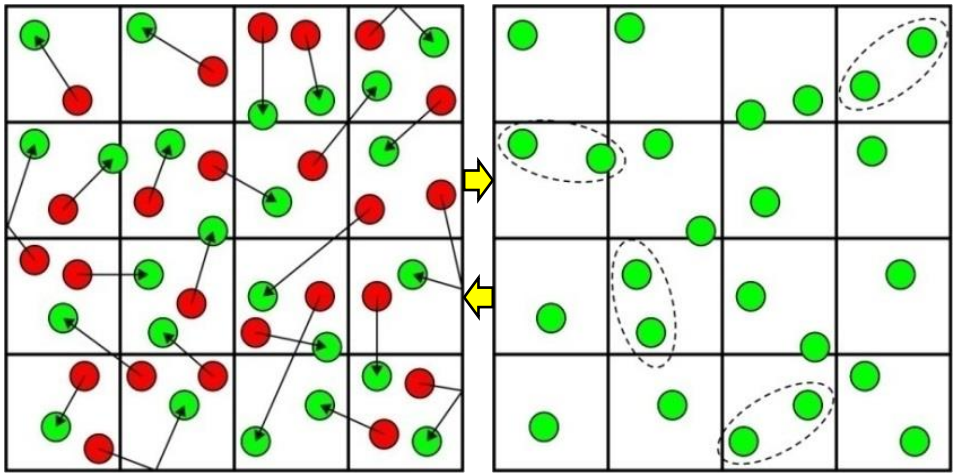

molecules move **molecules collide**

Four parameters control DSMC error:

Statistical error (1)

Samples per cell (S_c)

Discretization error (3)

- Particles per cell (N_c)
- Cell size (Δx)
- Time step (Δt)

Statistical and Particle-Number Errors

Sandia National Laboratories

Error related to sample size

- Statistical error
- Cell sample size $S_c = N_c \times N_t$
- N_c = particles per cell; N_t = time steps

Strategies for overcoming statistical error

- Use large number of samples
- For steady flows, use time and/or ensemble averaging
- Computational expense ~ S_c

Error related to local number of particles

- Error $\propto 1/N_c$
- Systematic – persists even as $S_c \rightarrow \infty$

Limited number of samples per time step

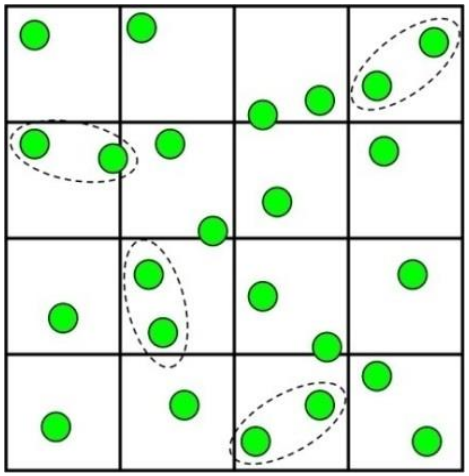

Not enough particles to capture physics

DSMC Convergence

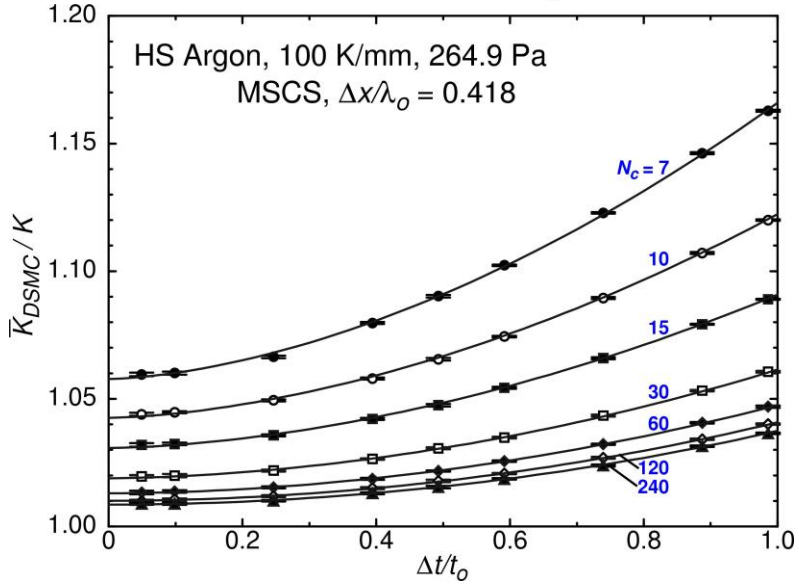

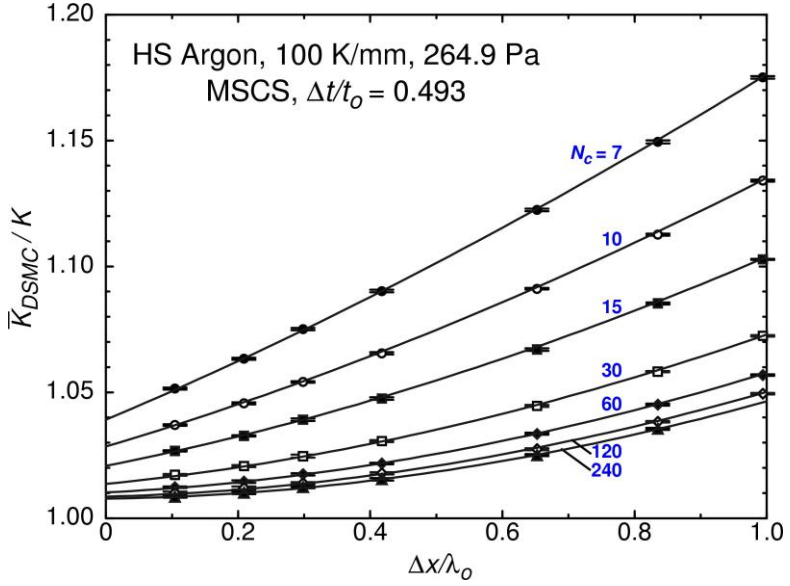

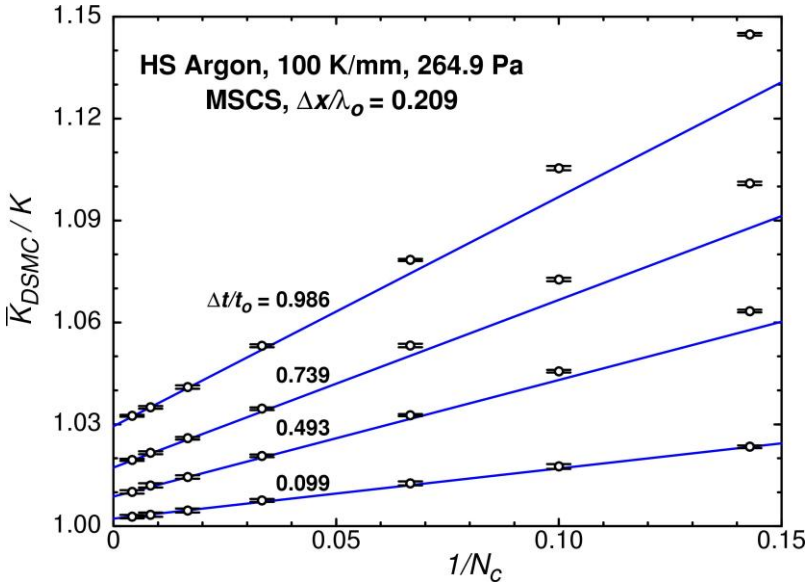

- Curves are best fits
- Error bars represent 95% confidence intervals
- Quadratic convergence for Δx, Δt
- First-order convergence $O(1/N_c)$, as $N_c \rightarrow \infty$
- Higher-order for long time steps
- For $N_c = 7$ and $\Delta t/t_o = 0.493$, convergence rate appears linear in $\Delta x/\lambda_o$

Functional Form of Error

Functional form that represents DSMC data

- *Ad hoc* series expansion in Δx, Δt, and $1/N_c$
- Perform least-squares fitting of entire data set

$$\frac{K_{DSMC}}{K} = 1.0000 + 0.0286\Delta\tilde{t}^2 + 0.0411\Delta\tilde{x}^2 - 0.0016\Delta\tilde{x}^3 - 0.023\Delta\tilde{t}^2\Delta\tilde{x}^2 +$$
$$-\frac{0.111}{N_c} + \frac{1}{N_c}\left[1.22\Delta\tilde{x} - 0.26\Delta\tilde{x}^2 + 0.97\Delta\tilde{t}^{3/2} + \ldots\right] + 0.95\frac{\Delta\tilde{t}^2}{N_c^2} + \ldots$$

Cross terms show convergence behavior is complex

Rader D. J., Gallis M. A., Torczynski J. R., Wagner W., "DSMC Convergence Behavior of the Hard-Sphere-Gas Thermal Conductivity for Fourier Heat Flow", *Phys. Fluids*, 18, 077102, 2006.

Infinite-Particle Convergence

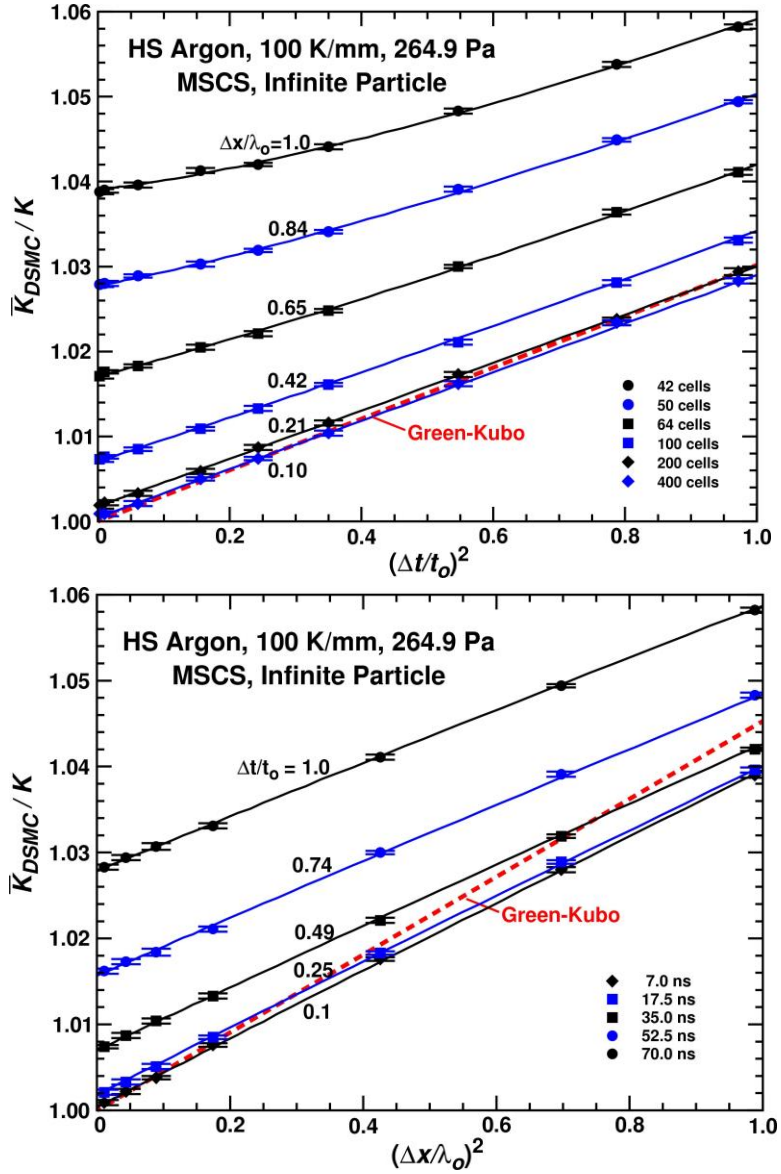

- Finite-particle error removed: values "extrapolated" to $N_c \rightarrow \infty$
- 63 extrapolated data points
- Error bars: fitting uncertainty
- Quadratic convergence in time step and cell size
- *Qualitative* agreement with *Green-Kubo* theory, but slopes are different
- Lines are best fits of data

Improving DSMC Nearest-Neighbor Procedures

Sandia National Laboratories

- Transient Adaptive Sub-Cells (Bird 2000)
  - Stochastically determines the nearest neighbor
  - Creates Cartesian sub-cell structure for each collision phase
  - Number of simulators in each sub-cell is ~ 1-2
  - Relatively inexpensive for large N: O(N)
- Virtual Sub-Cells (LeBeau 2003)
  - Deterministically determines the nearest neighbor
  - More accurate*
  - More expensive than TASC for large N: $O(N^2)$

DSMC07 = Sophisticated DSMC

DSMC94 = Established DSMC

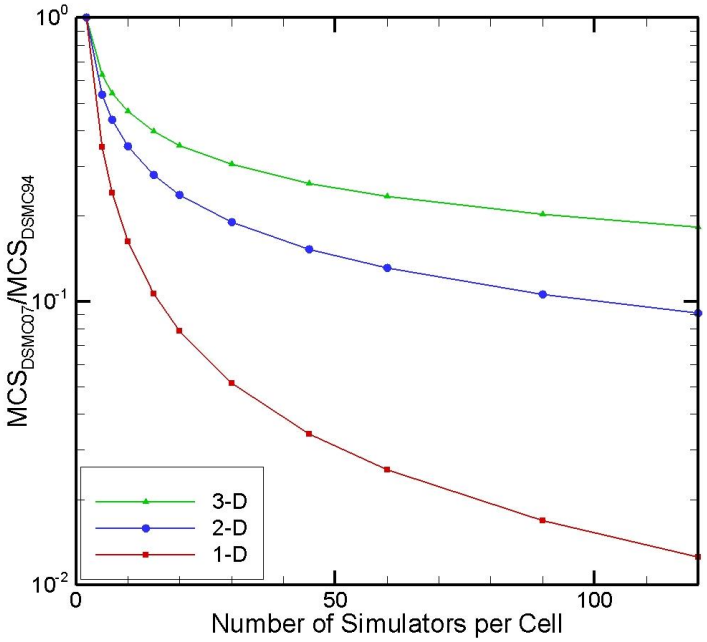

*A stress-tensor anisotropy is introduced in reduced-dimensionality simulations

Bird G. A., Gallis M. A., Torczynski J. R., Rader D. J., "Accuracy and efficiency of the sophisticated Direct Simulation Monte Carlo algorithm", Phys. Fluids, 21, 017103, 2009.

Hard-Sphere & Maxwell Transport Coefficient Profiles

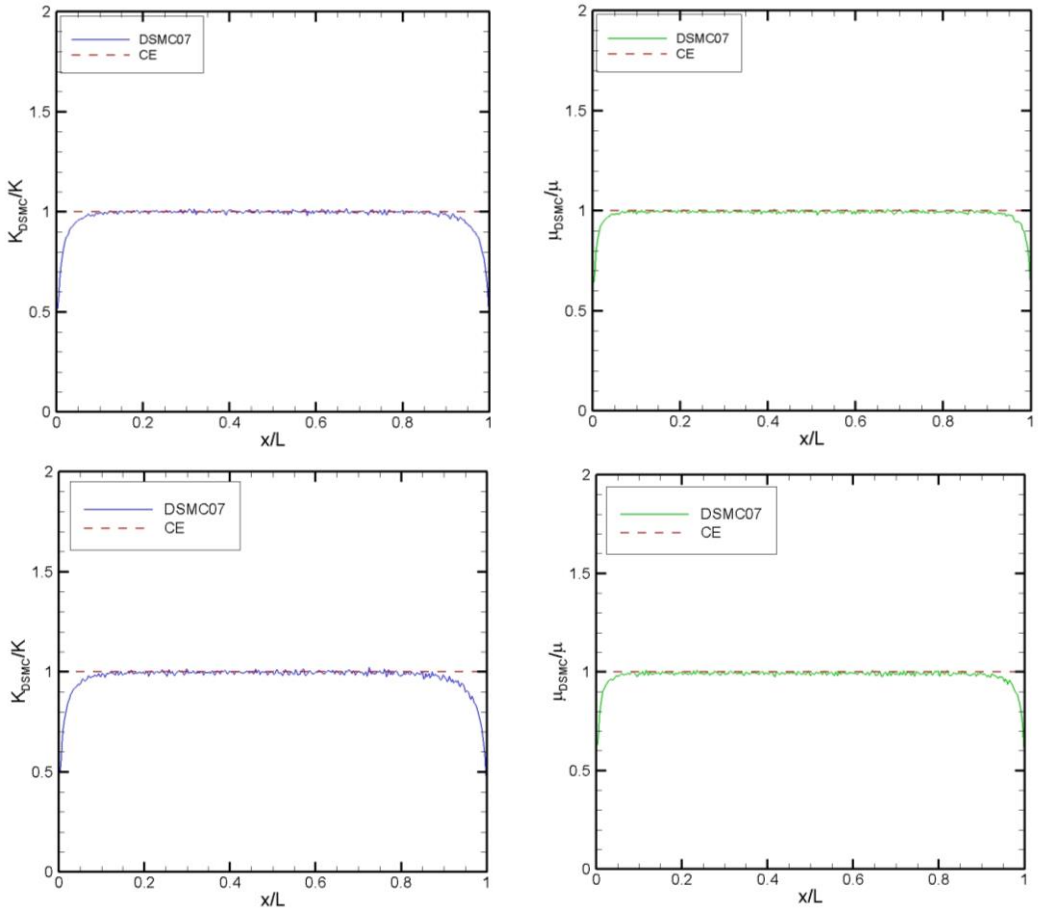

- DSMC07 and CE thermal conductivity and viscosity
  - Low heat flux, low shear stress: $Kn_q = 0.006$, $Kn_\tau = 0.003$
  - Agreement in central region: normal solution

Hard-Sphere & Maxwell Sonine-Coefficient Profiles

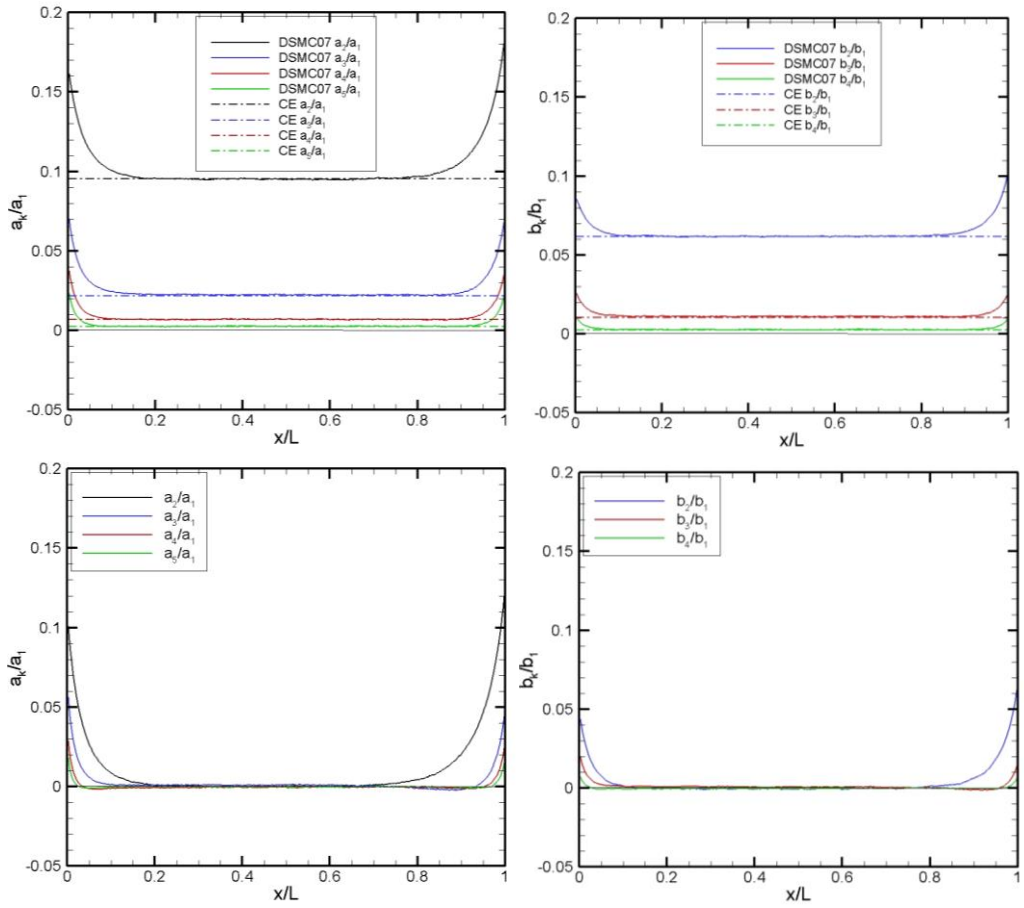

- DSMC07 and CE hard-sphere Sonine polynomial coefficients a_k/a_1 and b_k/b_1
  - Good agreement in central region: normal solution
  - Demonstrates accuracy of molecular velocity distribution

Maxwell Normal Sonine Coefficients

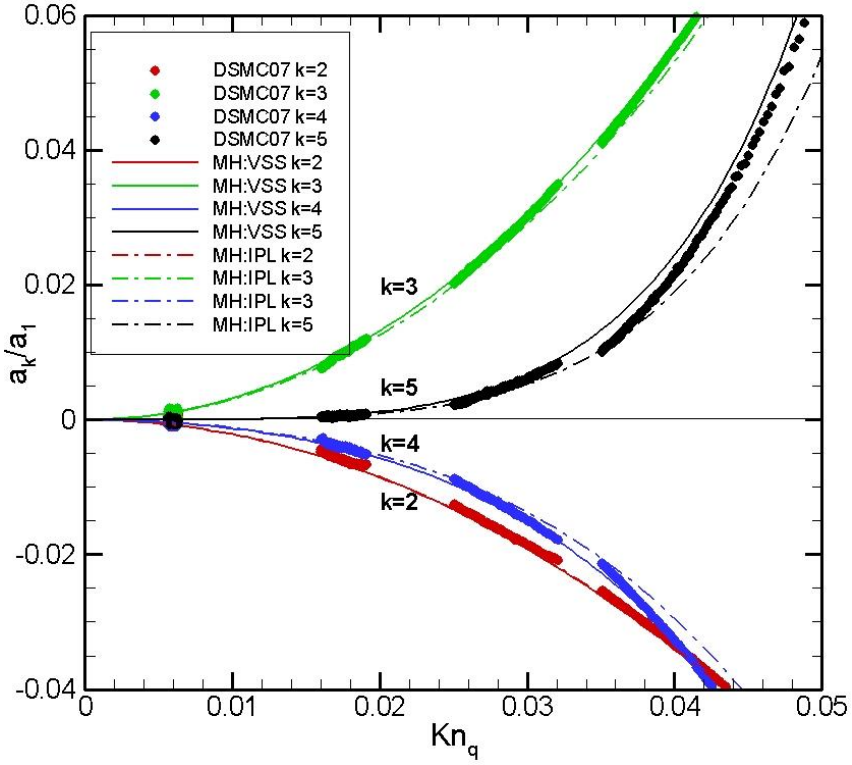

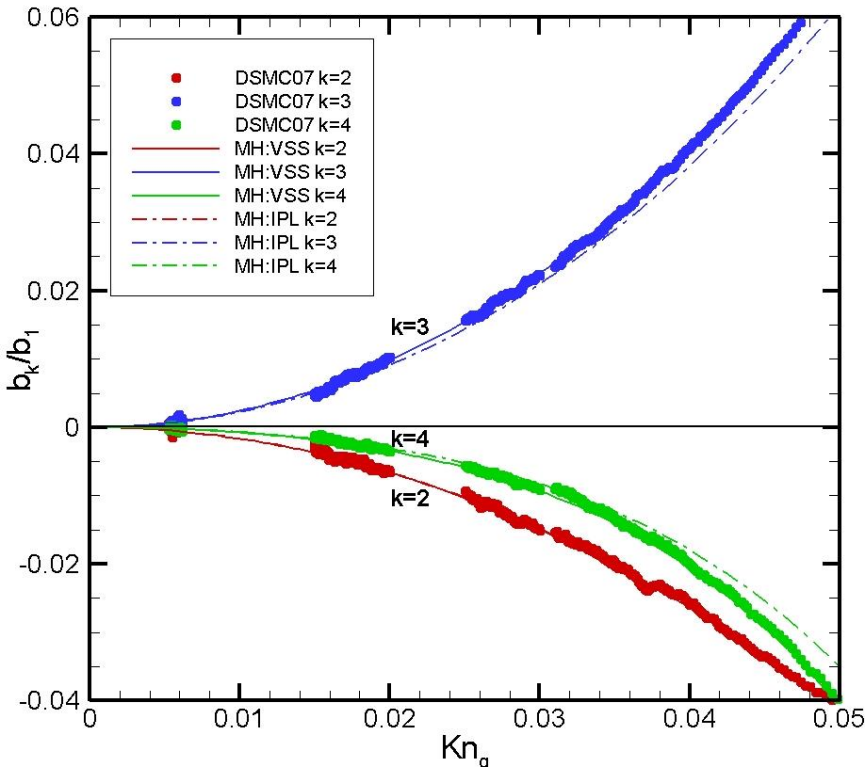

- DSMC07 and MH Maxwell normal solutions for a_k/a_1 and b_k/b_1
  - Four DSMC07 simulations: ΔT = 70, 200, 300, 400 K
  - MH: VSS-Maxwell (solid) and IPL-Maxwell (dashed) differ
  - DSMC07 and MH VSS-Maxwell normal solutions agree

DSMC-VSC Functional Form of Error

Sandia National Laboratories

Best-fit correlation function for sophisticated DSMC with $\Delta t = \Delta t_o$

$$\frac{K_{DSMC}}{K} = 0.9953 + 0.07773\left(\frac{2\cdot \mathrm{D}t}{t_o}\right) + 0.02795\left(\frac{2\cdot \mathrm{D}t}{t_o}\right)^2 - 0.0066\left(\frac{\mathrm{D}x}{l}\right)^2$$
$$- 0.0234\left(\frac{2\cdot \mathrm{D}t}{t_o}\right)^2\left(\frac{\mathrm{D}x}{l}\right)^2 + 0.60375\frac{1}{N_c} + \frac{1}{N_c}F\left[\frac{2\cdot \mathrm{D}t}{t_o}, \frac{\mathrm{D}x}{l}, \left(\frac{2\cdot \mathrm{D}t}{t_o}\right)^2\right]$$

DSMC limiting convergence differs from Green-Kubo (GK) behavior

- Weak quadratic convergence in cell size ($\Delta t/t_o \to 0$, $N_c \to \infty$)
- Linear convergence in $1/N_c$ for $N_c \geq 30$ simulators/cell
- Linear convergence in time step ($\Delta x/\lambda \to 0$, $N_c \to \infty$)

Gallis M. A., Torczynski J. R., Rader D. J., Bird G. A, "Convergence behavior of a new DSMC algorithm", J. Comput. Phys. 228 , 4532-4548, 2009.

Convergence Behavior for $N_c = 10$ Effect of Cell Size

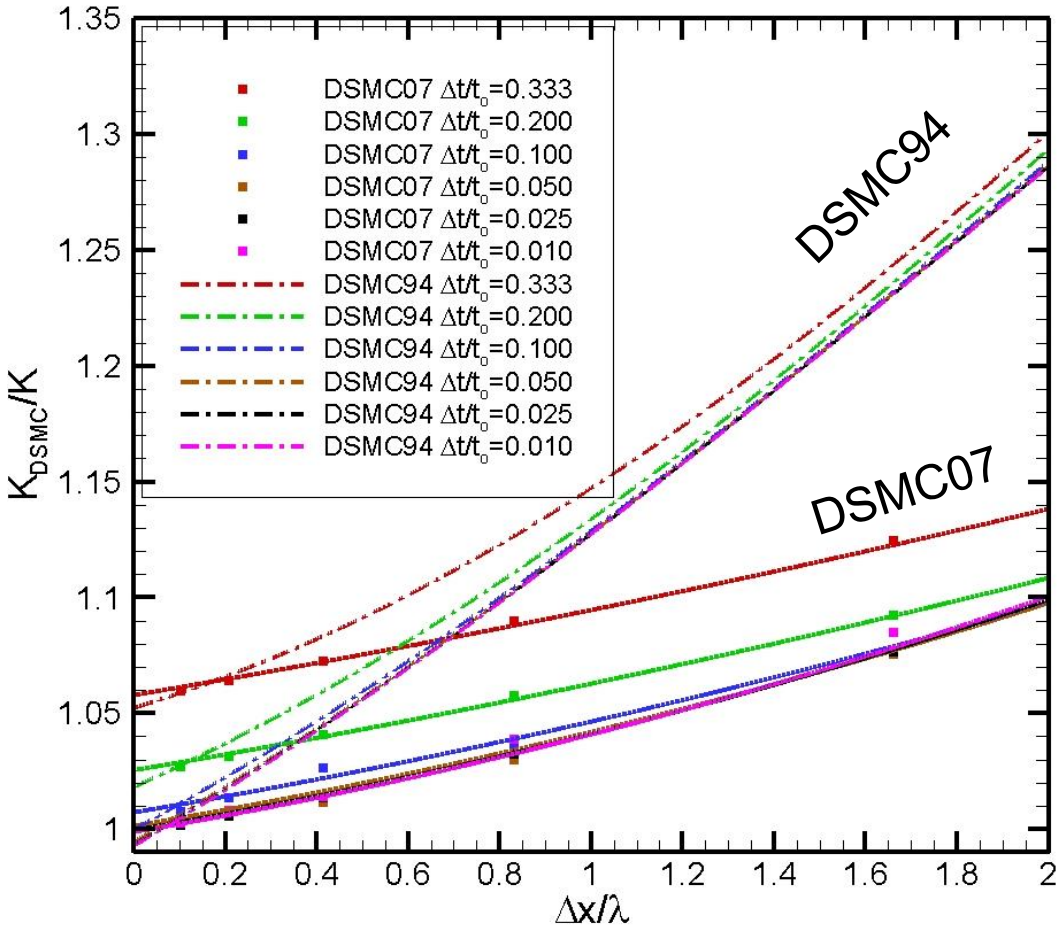

- For finite number of simulators, the algorithm
  - Is insensitive to spatial resolution
  - Has error below 2% for all cases as long as $\Delta t/t_o \leq 0.2$

Convergence Behavior for N_c = 15 & 30 Effect of Cell Size

Sandia National Laboratories

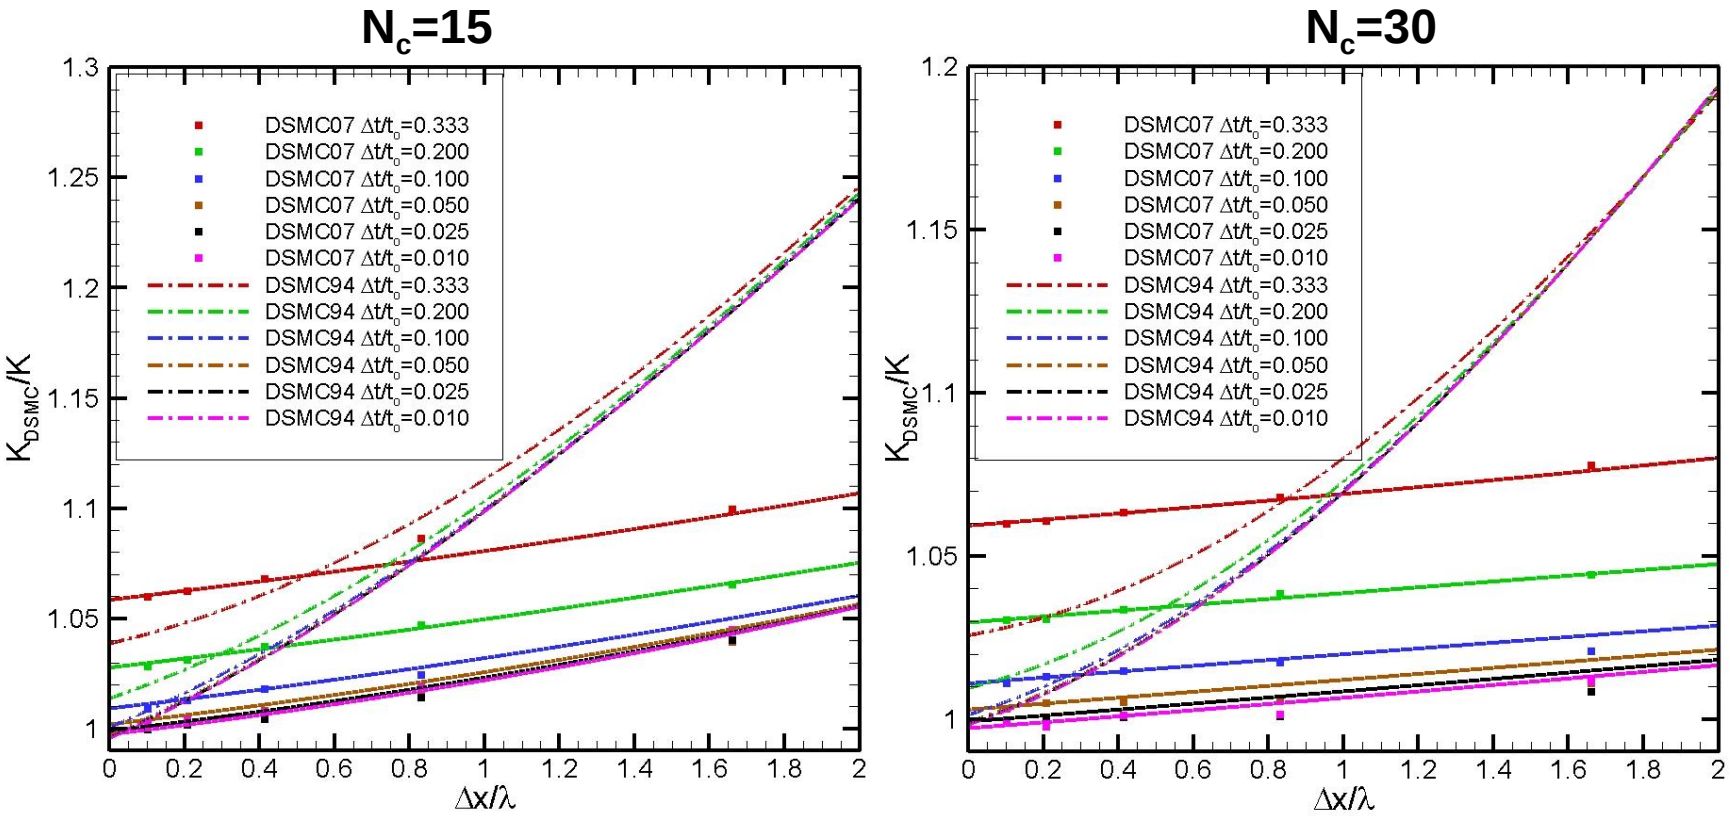

For a finite number of simulators, the algorithm

- Is insensitive to spatial resolution.

DSMC07 error is lower than DSMC94 only when

- $\Delta t/t_o \leq 0.1$ (Nc=15), $\Delta t/t_o \leq 0.05$ (Nc=30).

Convergence Behavior for N_c = 30 Selection Limit 10 & No Limit

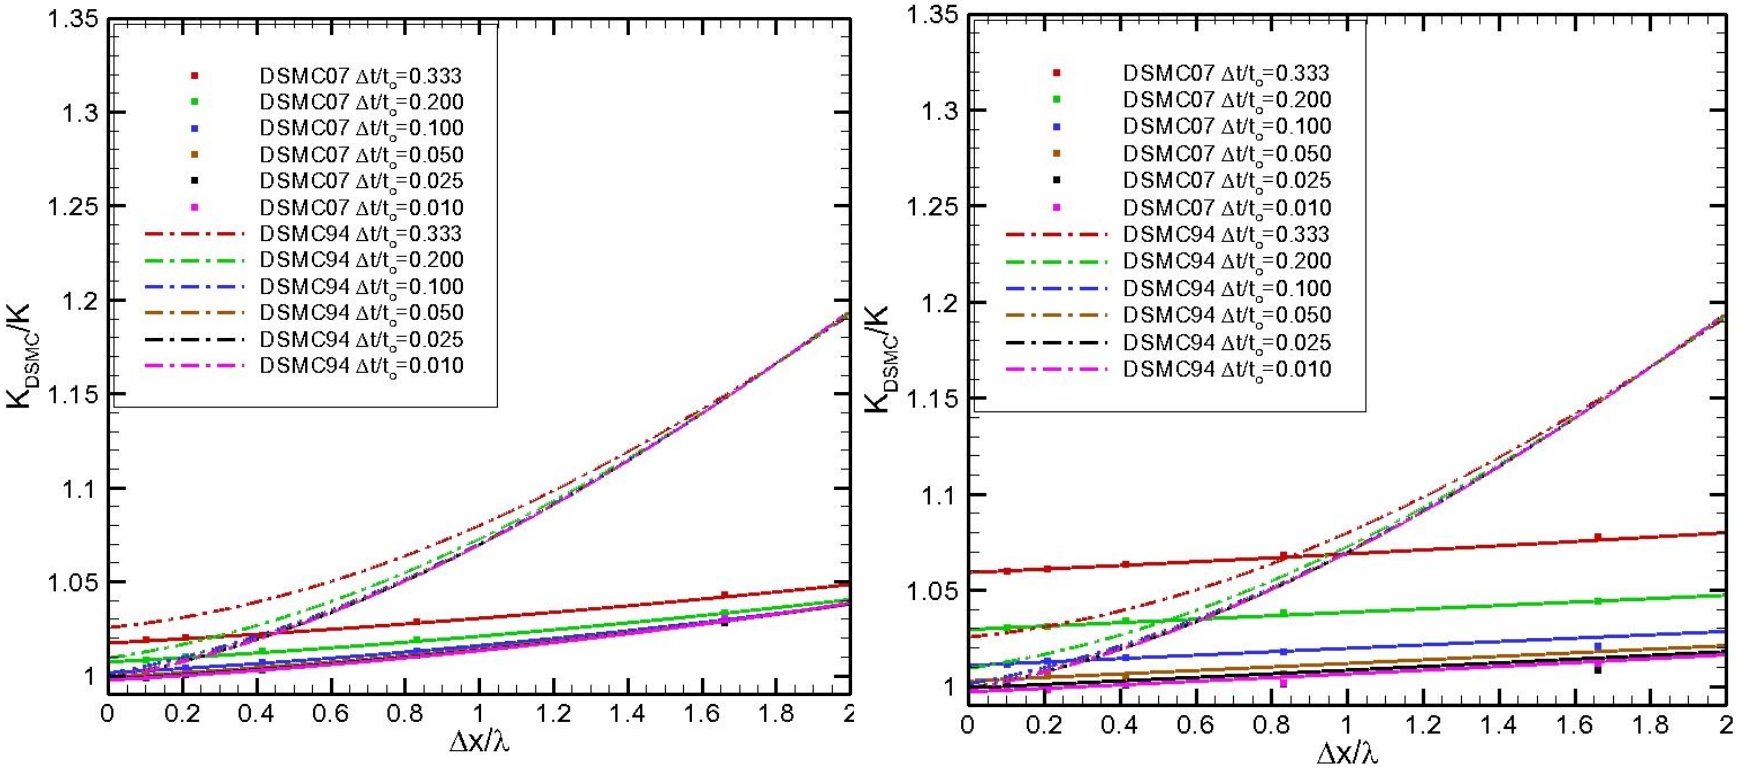

- Significant performance improvement with smaller computational cost can be achieved when a random subset of simulators is used.
- A limit of 30 should be adequate for most applications, as suggested by Bird.

Convergence Behavior for $N_c = 15$ Trajectory dependent selection

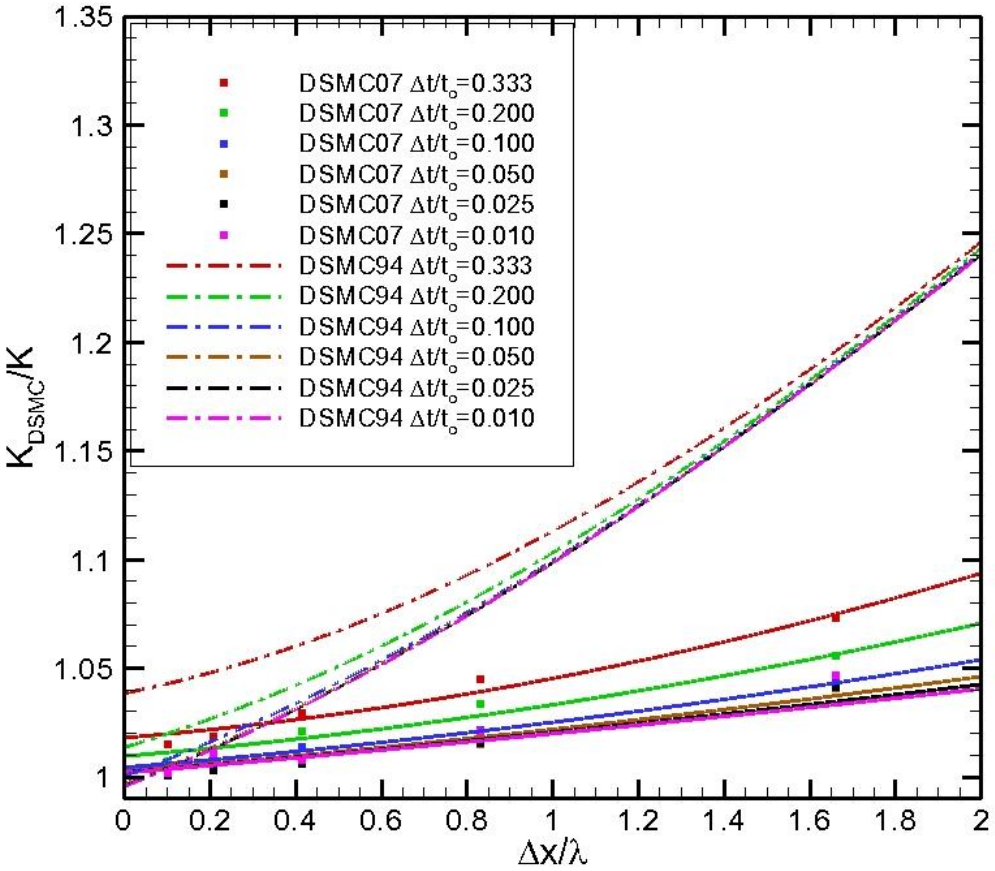

Simulators with high velocities have more simulators available as collision partners, whereas simulators with low velocities have to collide with their nearest neighbor.

To minimize the time-step error in the VSC scheme, a near-neighbor collision partner is selected from within a sphere centered on the simulator with a radius proportional to the distance traveled by the simulator during the latest advection phase.

Gallis, M. A. Torczynski J. R., "Effect of collision-partner selection schemes on the accuracy and efficiency of the direct simulation Monte Carlo method", International Journal for Numerical Methods in Fluids, Vol. 67, No. 8, pp.1057-1072, 2011.

Parallel Efficiency: The Unfair Advantage

Sandia National Laboratories

- The advantages of DSMC come at a cost
- DSMC is computationally efficient but computationally intense
- Its successful application to real problems depends heavily on its parallel performance

- 1000x speedup required for some problems of interest
- Monte Carlo methods usually have good parallel performance
  - The workload depends mainly on the simulators within a cell
  - Relatively less need to communicate information between cells
  - Trivial to parallelize in velocity space

The necessary speedup can be achieved without any loss of accuracy or convergence characteristics through parallel computing

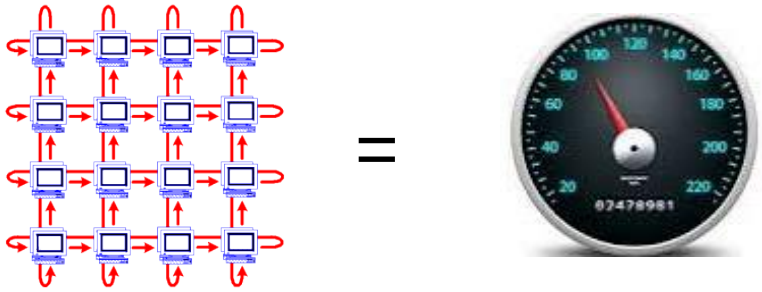

Top 5 Supercomputers (2014)

Sandia National Laboratories

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)
1	National Super Computer Center in Guangzhou	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P	3,120,000	33,862.7	54,902.4
2	DOE/SC/Oak Ridge National Laboratory	**Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x	560,640	17,590.0	27,112.5
3	DOE/NNSA/LLNL	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom	1,572,864	17,173.2	20,132.7
4	RIKEN Advanced Institute for Computational Science (AICS)	**K computer**, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	705,024	10,510.0	11,280.4
5	DOE/SC/Argonne National Laboratory	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom	786,432	8,586.6	10,066.3

Programming for Next Generation and Exascale Machines

Sandia National Laboratories

- Millions of nodes likely
- Reduced memory per node
- Parallelism within node:
  - Multi-core: 16 and growing
  - Many-core: Intel Xeon Phi, 240 threads
  - GPUs: NVIDIA/AMD, 1000 warps
- Example:
  - LLNL BG/Q: 96K nodes, 16 cores/node + 4 MPI tasks/core

Programming model: MPI + X

- Goal is to decouple the science code from the hardware details

Necessary elements

- Adaptive gridding
- In-situ visualization
- Efficient communications
- Load balancing

Aiming for MPI+X via Kokkos

Sandia National Laboratories

- What is Kokkos:
  - Programming model in development at Sandia
  - C++ template library
  - Open-source
  - Stand-alone
- Goal: write application kernels only once, and run them efficiently on a wide variety of hardware platforms
- Two major components:
  - Data access abstraction via Kokkos arrays optimal layout & access pattern for each device: GPU, Xeon Phi, etc.
  - Parallel dispatch of small chunks of work auto-mapped onto back-end languages: CUDA, OpenMP, etc.

Developing an Exascale DSMC Code

Sandia National Laboratories

SPARTA = Stochastic PArallel Rarefied-gas Time-accurate Analyzer

General features

- 2D or 3D, serial or parallel
- Cartesian, hierarchical grid
  - Oct-tree (up to 16 levels in 64-bit cell ID)
  - Multilevel, general NxMxL instead of 2x2x2
- Triangulated surfaces cut/split the grid cells
  - 3D via Schwartzentruber algorithm
  - 2D via Weiler/Atherton algorithm
  - Formulated so can use as kernel in 3D algorithm
- C++, but really object-oriented C
  - Designed to be easy to extend
  - New collision/chemistry models, boundary conditions, etc.

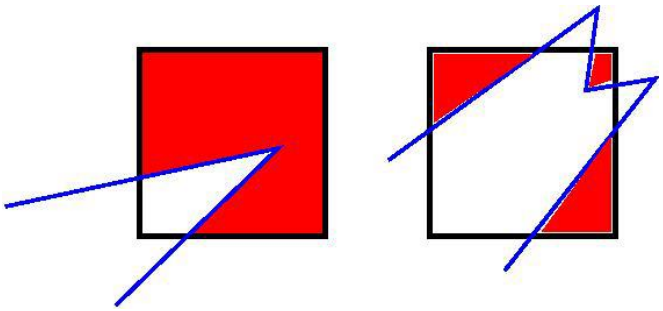

Adaptive Gridding

Sandia National Laboratories

- Create/adapt grid in situ, rather than pre-process & read in
- Examples: Generate around surface to user-specified resolution, adapt grid based on flow properties
- Algorithms should be efficient if they require only local communications

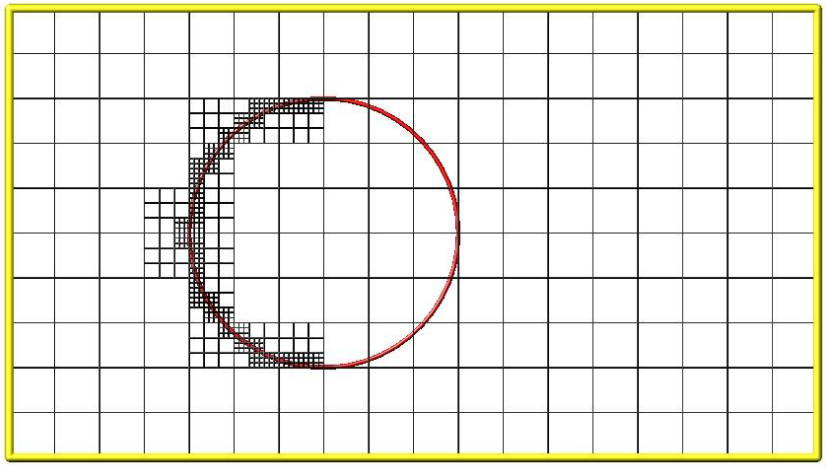

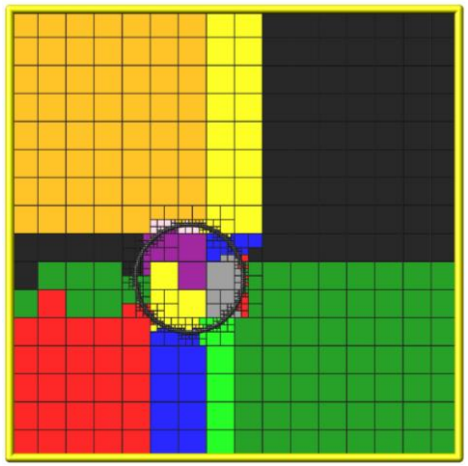

- Another setup task: label cells as outside/inside
- Simple if pre-processing, in situ easier for large problems

Simulation of Complicated Shapes

Sandia National Laboratories

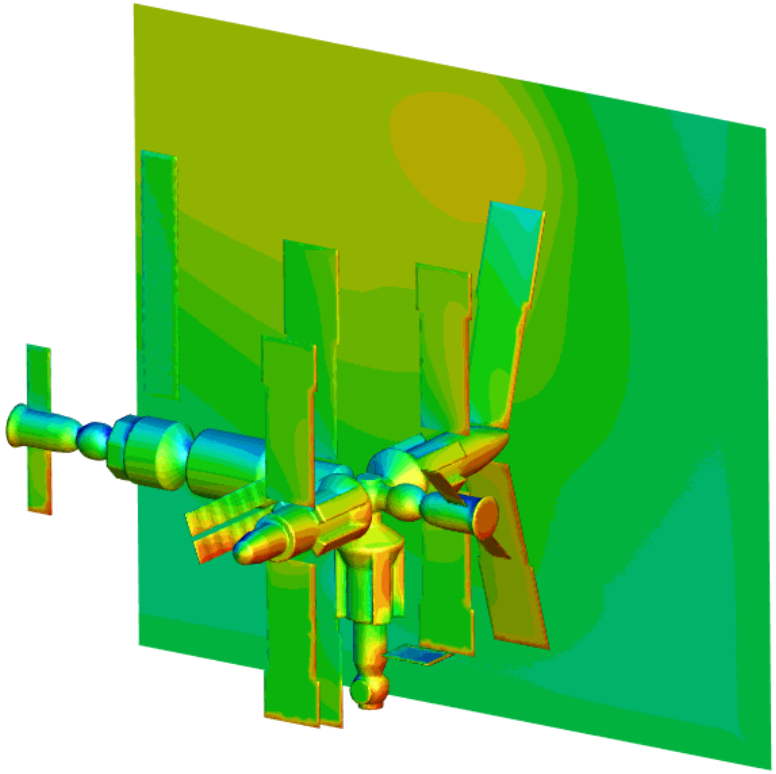

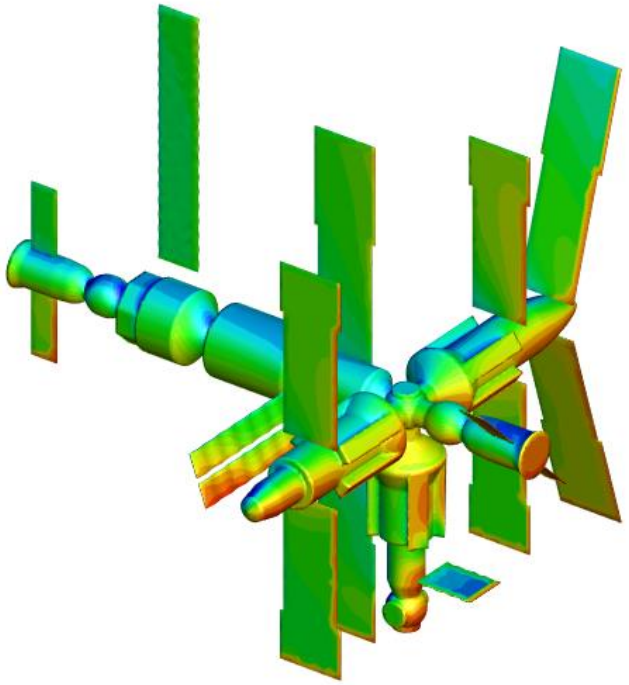

Grid generation (10^7 cells) completed in 0.3 seconds on 16 processors
Geometry comprises multiple “water-tight” bodies

In-Situ Visualization

Sandia National Laboratories

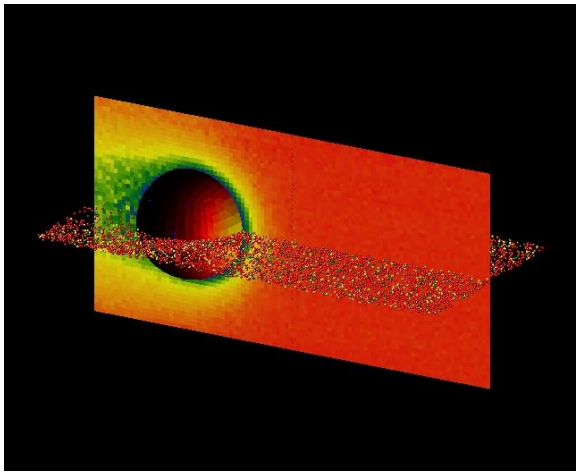

Not a replacement for interactive viz, but ...
Quite useful for debugging & quick analysis
At end of simulation (or during), instant movie

Render a JPG snapshot every N time steps:

- Each processor starts with blank image (1024x1024)
- Processor draws its cells/surfaces/molecules with depth-per-pixel
- Merge pairs of images, keep the pixel in front, recurse
- Draw is parallel, merge is logarithmic (like MPI Allreduce)

Images are ray-traced quality

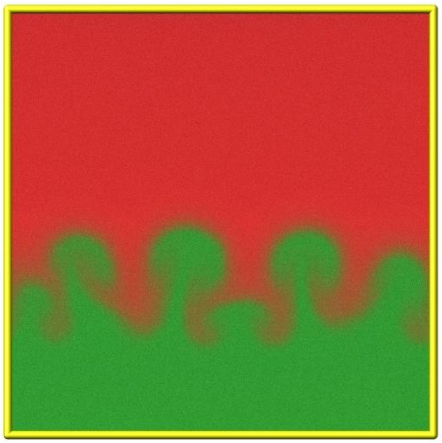

Load Balancing

Sandia National Laboratories

Balance across processors, static or dynamic
Granularity = grid cell with its molecules
Geometric method: recursive coordinate bisection (RCB)
Weighted by cell count or molecules or CPU

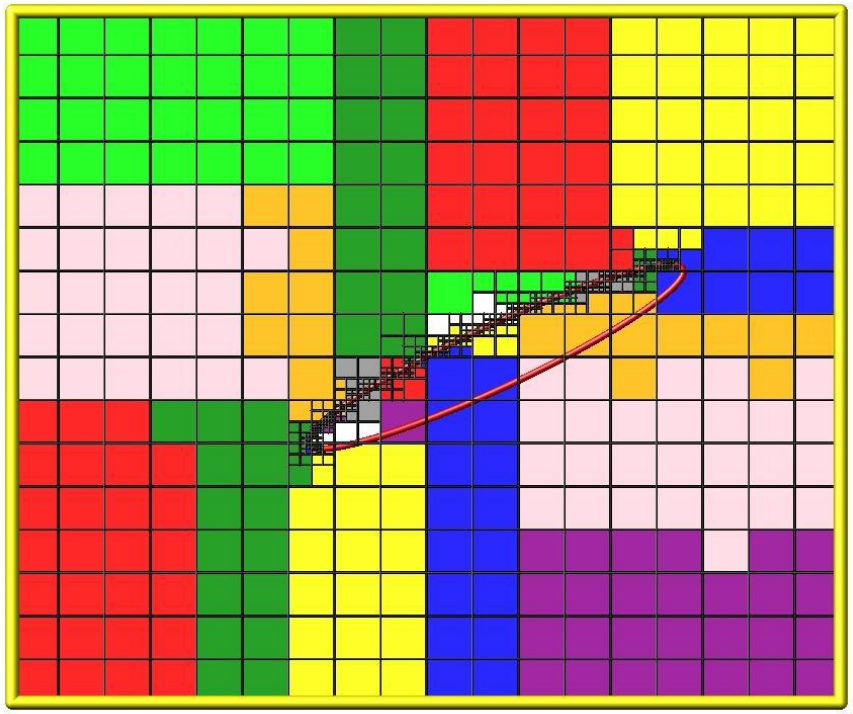

RCB is fast
Bigger cost is data move
Example:
1B cells on1024 BG/Q nodes
Worst case: move all cells
Balance time = 15 s:
(RCB=2, move=12, ghosts=1)

Efficient Communication

Sandia National Laboratories

- One processor = compact clump of cells via load balancing
  - Ghost region = nearby cells within user-defined cutoff
  - Store surface information for ghost cells to complete move

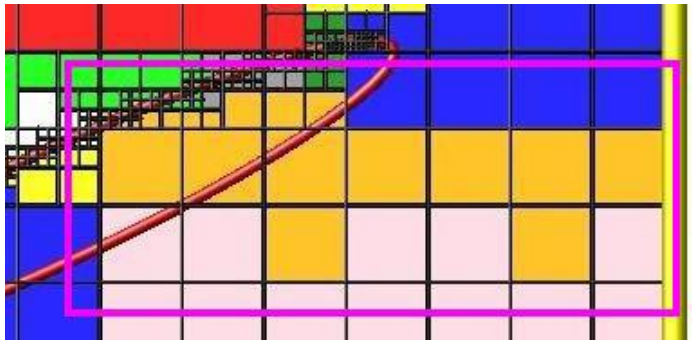

- Efficiently distributes grid information across processors
  - With sufficient cutoff, only one communication per step
  - Multiple passes if needed (or can bound molecule move)
- Communication with modest count of neighbor processors

SPARTA Benchmarking

Sandia National Laboratories

2 test cases:

- Free-molecular
  - Stress test for communication
  - 3D regular grid, 10^4-10^{11} (0.1 trillion) grid cells
  - 10 molecules/cell, 10^5-10^{12} (1 trillion) molecules
- Collisional
  - About 2x slower (sorting, collisions)
  - Same grid cell & molecule counts
- Effect of threading
  - 4 threads/core = 2x speed

SPARTA Benchmarking

16 cores/node
1 task/core

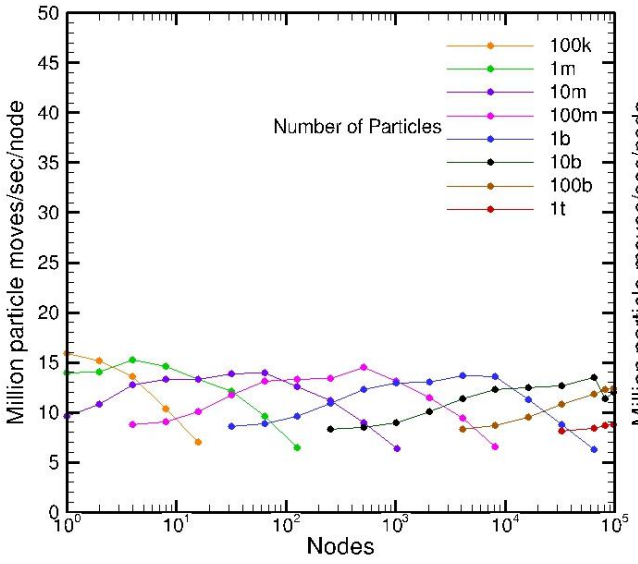

16 cores/node
2 tasks/core

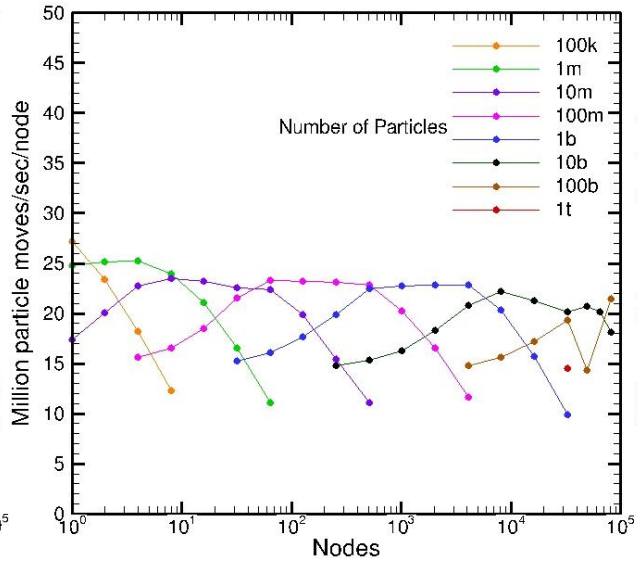

16 cores/node
4 tasks/core

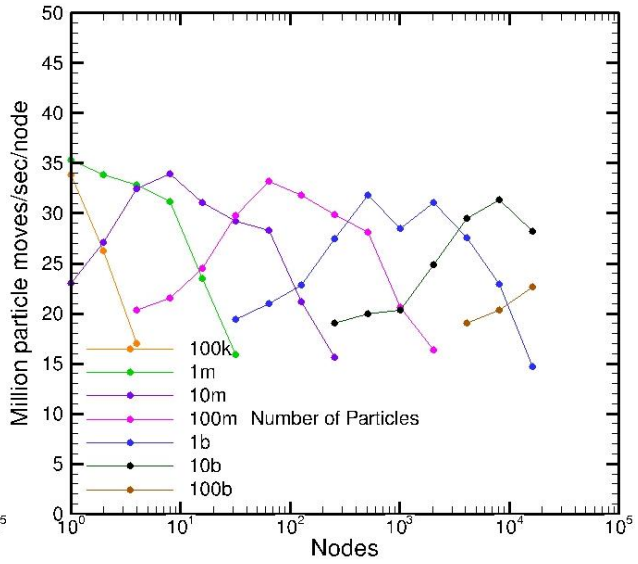

- Weak scaling indicates, 10% peak performance reduction from 1 to 10^6 cores
- 2 tasks/core gives 1.5x speedup, 4 tasks/core gives 2x speedup
- A total of 1 trillion simulators can be simulated on one third of the BG/Q
- Maximum number of tasks is 2.6 million

SPARTA Benchmarking (FM)

16 cores/node, 1 task/core

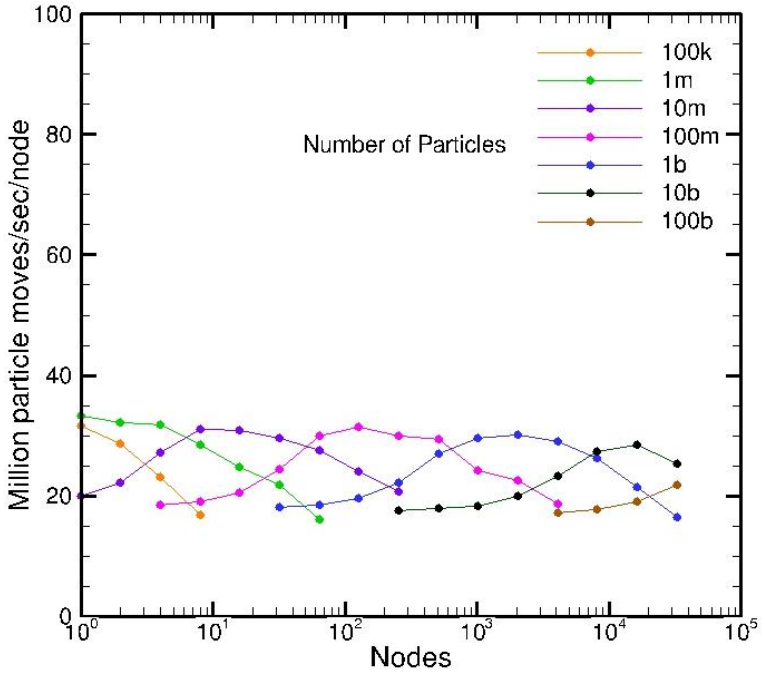

16 cores/node, 4 tasks/core

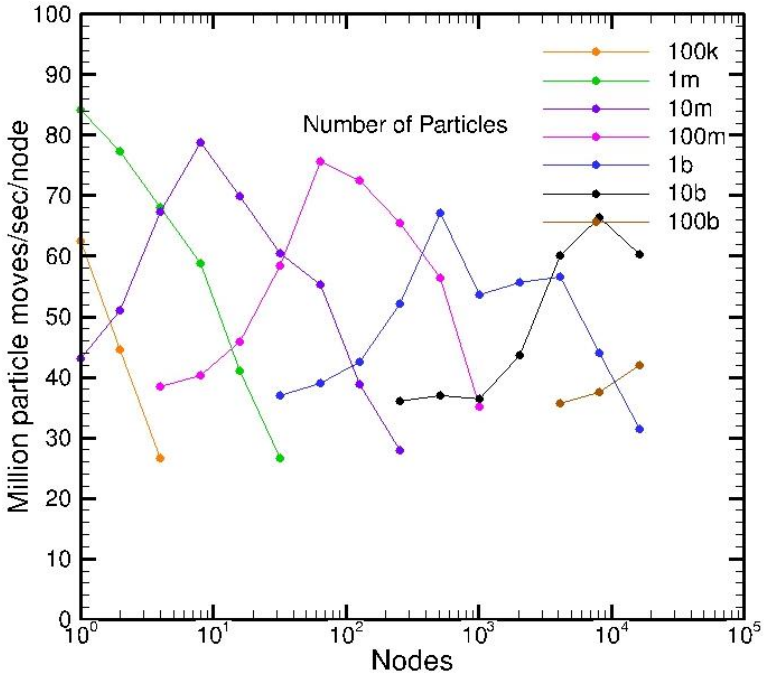

- Free-molecular (FM) calculations stress communications
- 2x speedup compared to collisional

Richtmyer-Meshkov Instability (RMI)

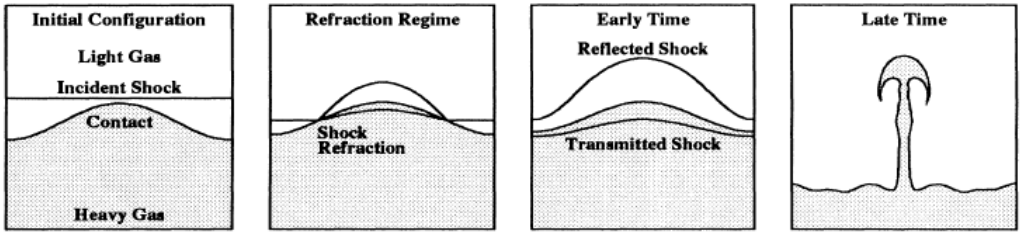

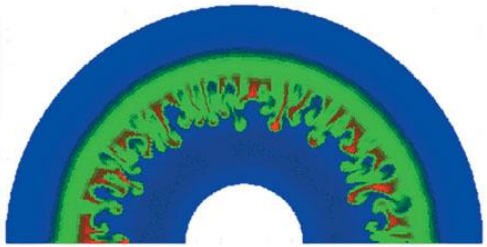

ICF target compression

Applications include Inertial Confinement Fusion (ICF), stellar evolution models, interaction of shocks with flames

RMI combines multiple compressible phenomena

- Shock interaction, refraction, reflection, transmission
- Hydrodynamic instability, including:
- Nonlinear growth
- Subsequent transition to turbulence
- Range of Mach numbers
- Chemical reactions (combustion)

RMI in He/Ar Mixture: Mach 1.2 Shock

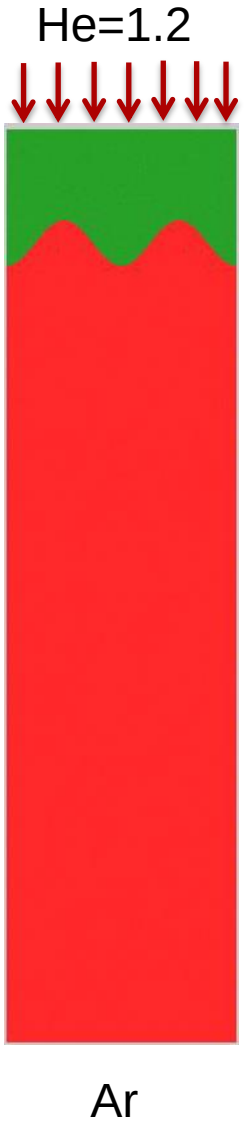

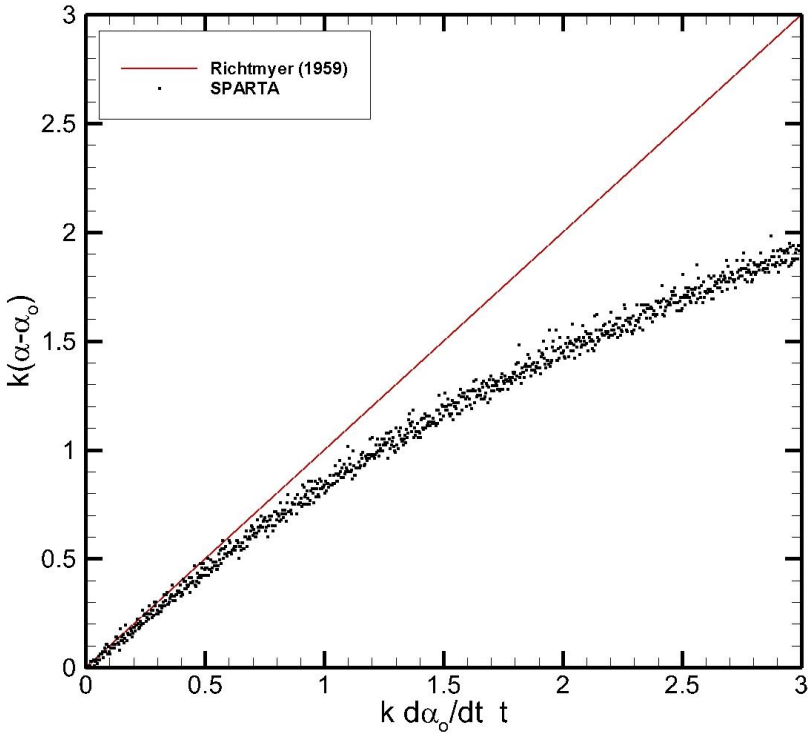

Non-dimensional amplitude for an initially small amplitude perturbation compared to Richtmyer's model for early time evolution

RMI in Air-SF_6 Mixture: Mach = 1.2 Shock

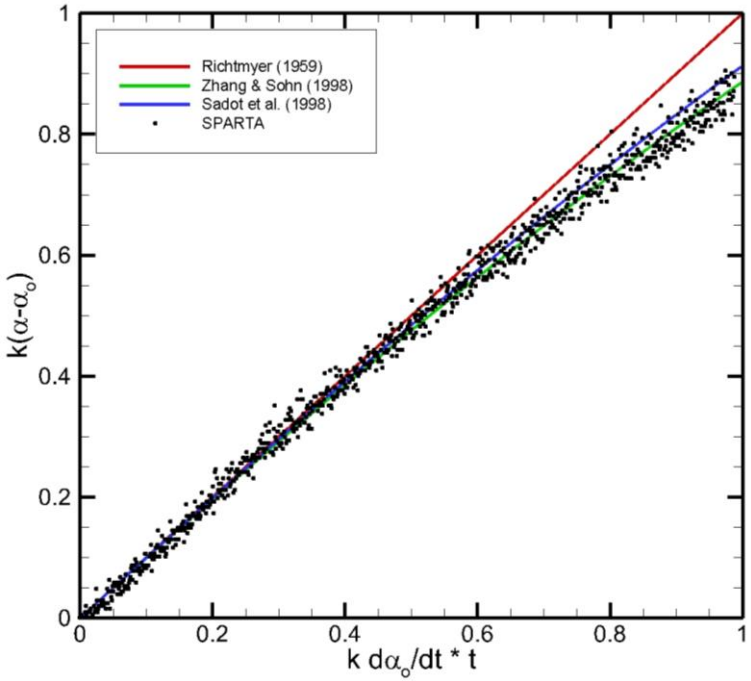

Non-dimensional amplitude for an initially small amplitude perturbation compared to theoretical/empirical models

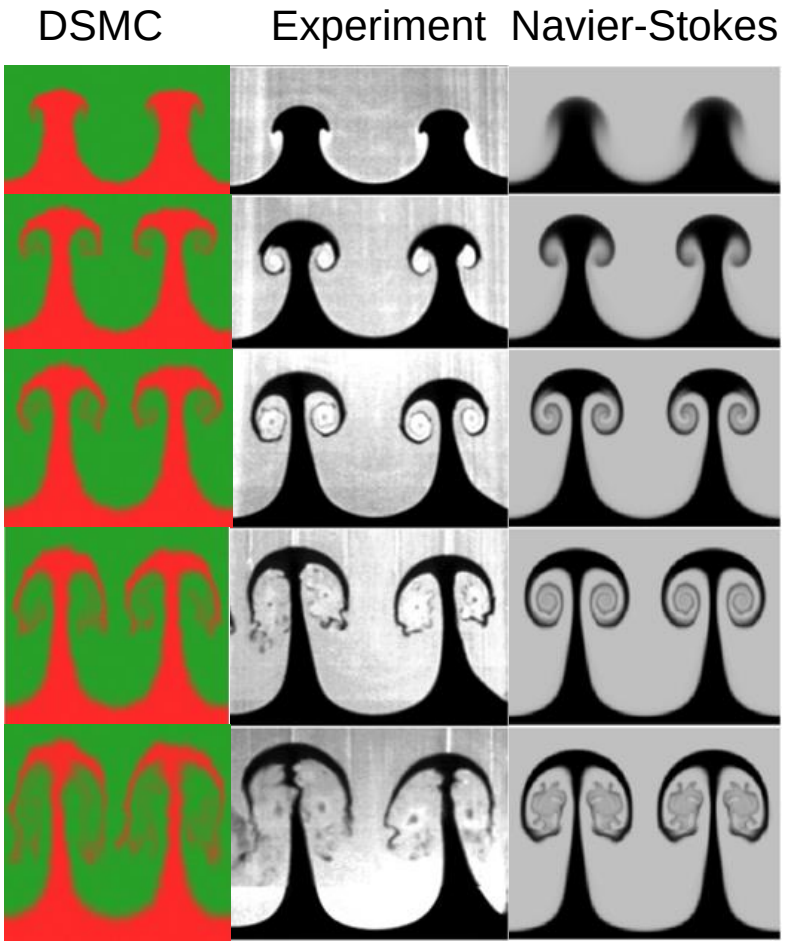

Conclusions

DSMC yields exquisite agreement with analytical results, where available

- Chapman-Enskog, Moment-Hierarchy theory
- Discretization & sampling errors understood

DSMC scales extremely well & can take full advantage of massively parallel platforms

- Can simulate unprecedented flow regimes
- Hydrodynamic instabilities, lower altitudes

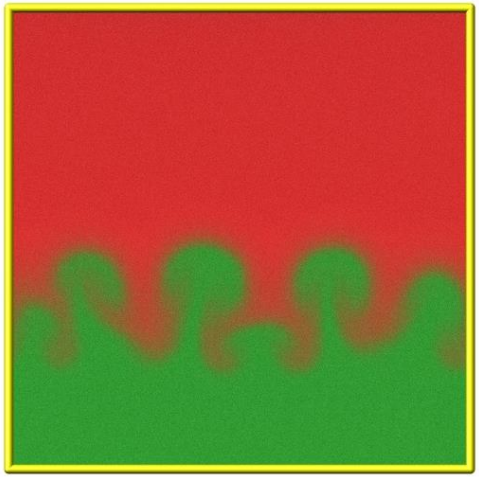

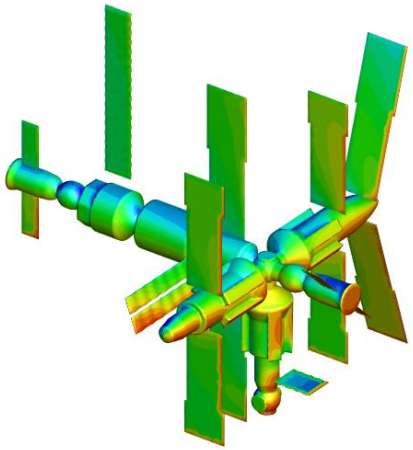

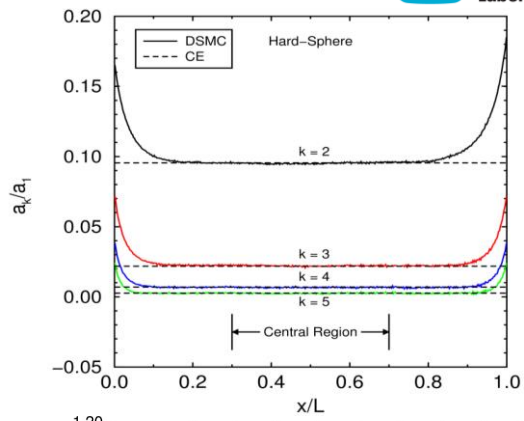

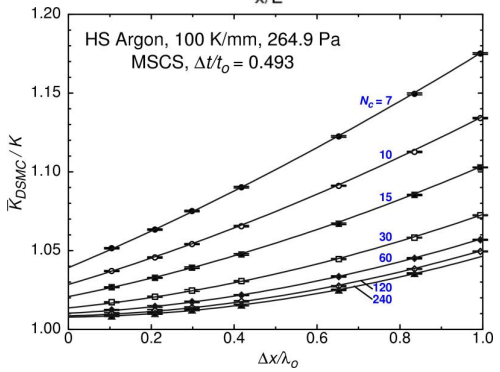

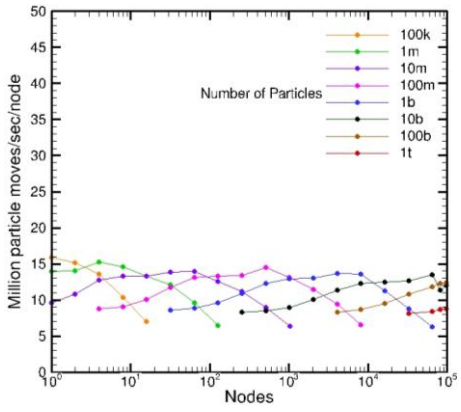

Thank you!

Sandia National Laboratories

Graeme Bird

DSMC15

Sandia National Laboratories

DSMC15 Workshop, Hawaii, September 2015

Be there.
Aloha!