SAND2015-3134PE

Exceptional service in the national interest

Sandia National Laboratories

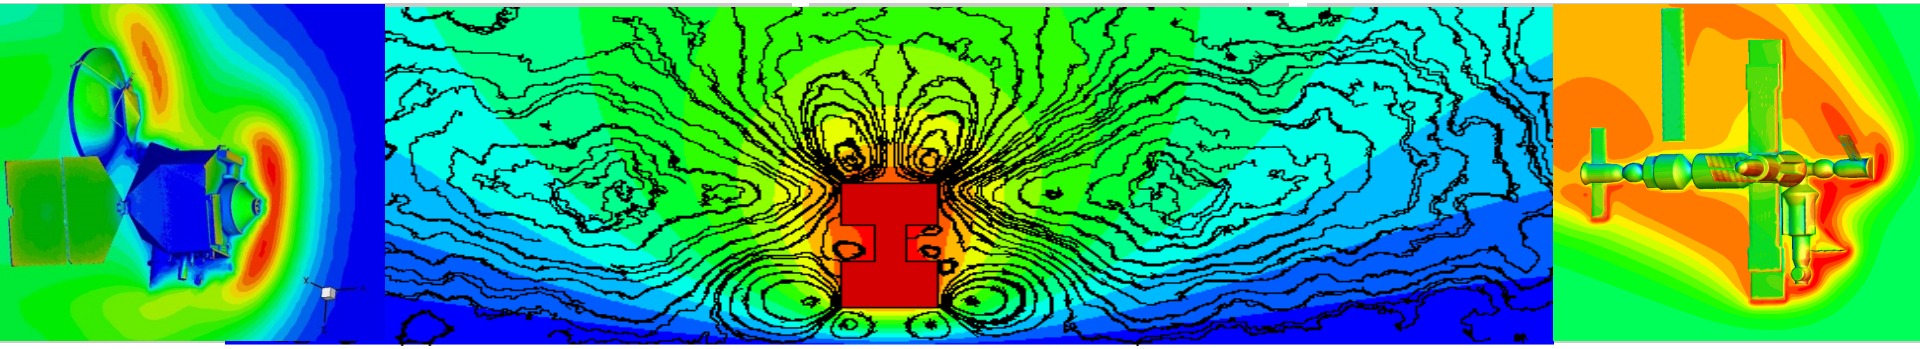

Stochastic PArallel Rarefied-gas Time-accurate Analyzer

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.

Length Scales for Dilute Gases

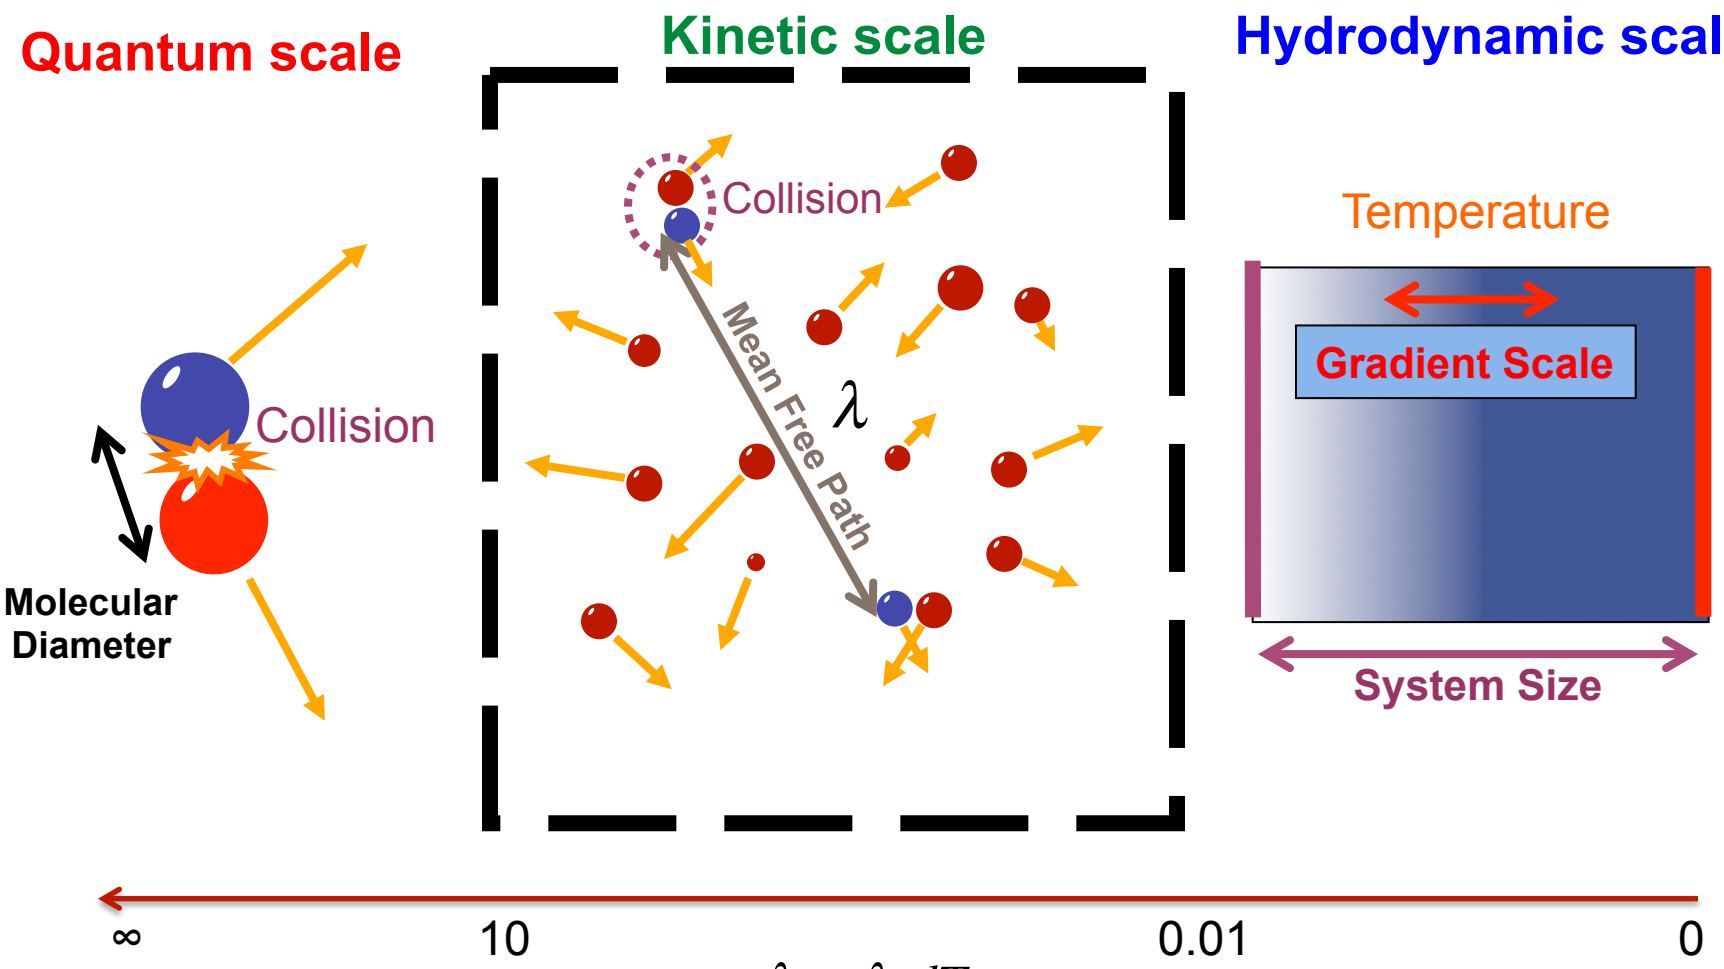

$$Kn = \frac{\lambda}{L} = \frac{\lambda}{T}\frac{dT}{dx}$$

Direct Simulation Monte Carlo

Sandia National Laboratories

How DSMC works

DSMC molecule-simulators **statistically** represent a large number of real molecules ($O(10^{10})$-$O(10^{15})$)

Computational molecules move ballistically, collide statistically, and interact statistically with surfaces **like real molecules**

Molecular movement, surface-interaction, and collision are implemented **sequentially** in the algorithm

Cell-based molecular statistics ("moments") are sampled and averaged over many time steps for steady flow

DSMC is inherently a transient method

Steady state is the ensemble average of unsteady state moves

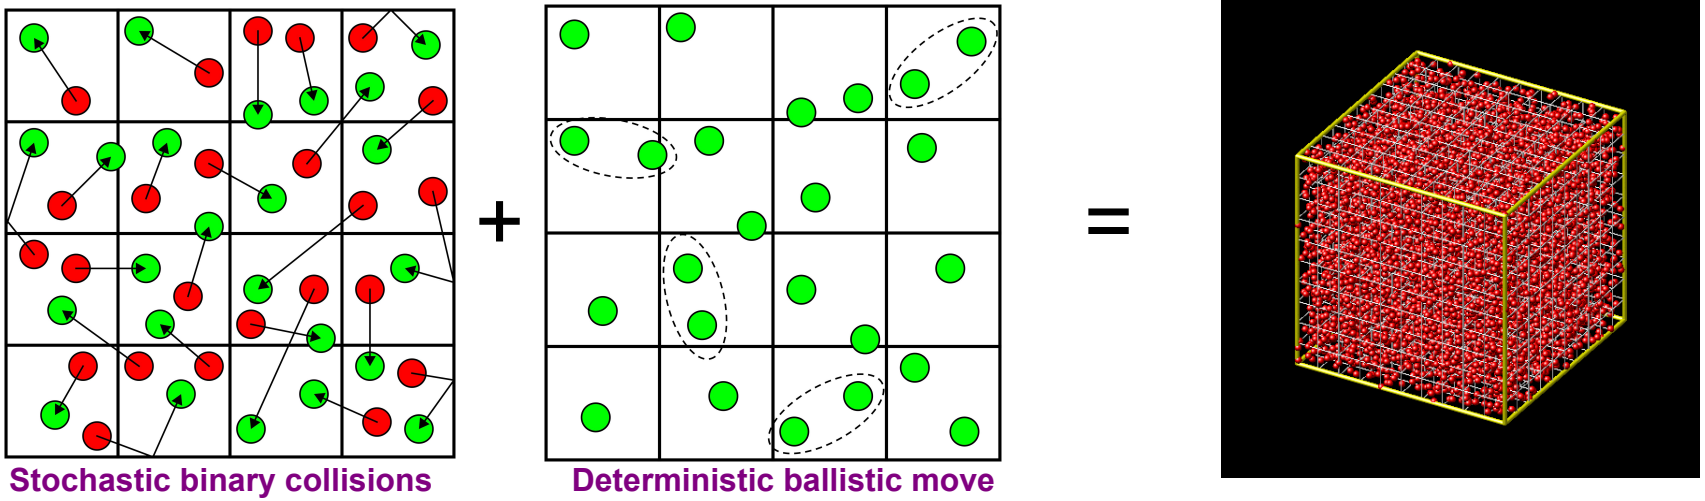

Developing an Exascale DSMC Code

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.

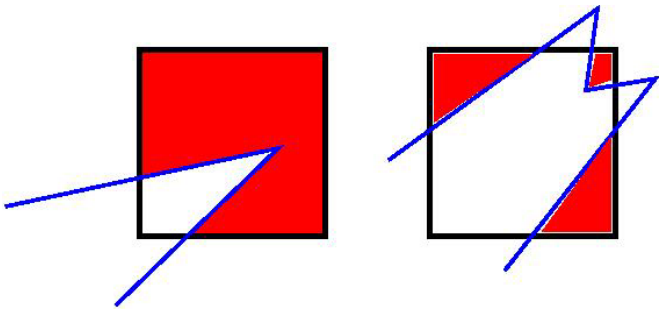

- Code available at http://sparta.sandia.gov

Earth Return Vehicles

Sandia National Laboratories

Gemini 1961-1966	Apollo 1963-1972	Space Shuttle 1981-2011	Orion 2014-?
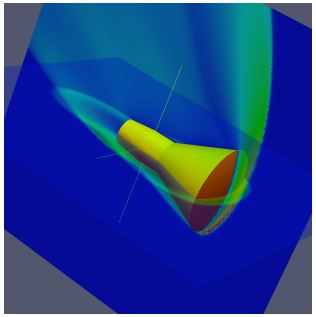	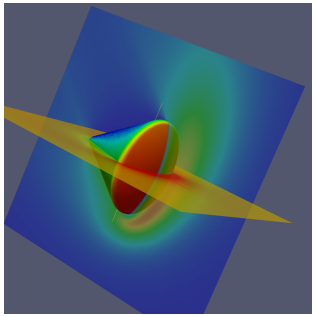	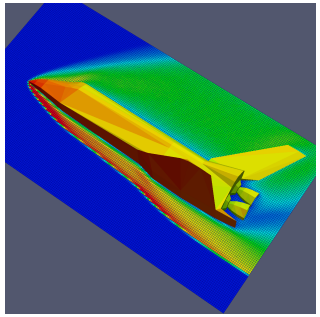	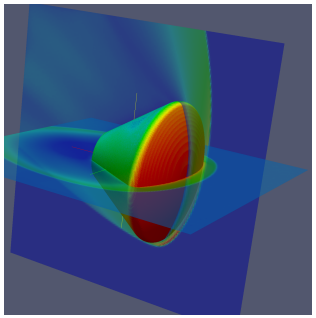
			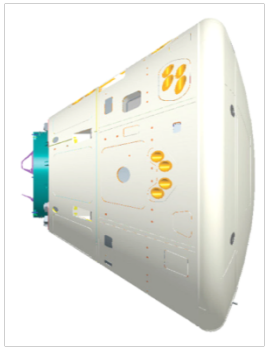

- Atmospheric entry system must provide controlled dissipation of kinetic and potential energy of the vehicle.
- Dynamic and thermal loads must be kept with certain range

Simulation of Complicated Shapes

Sandia National Laboratories

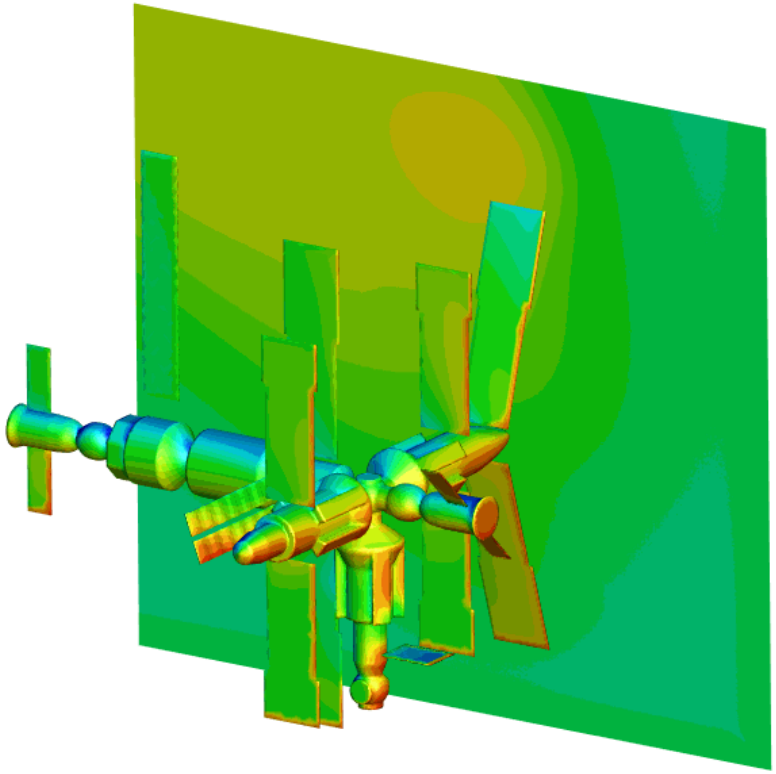

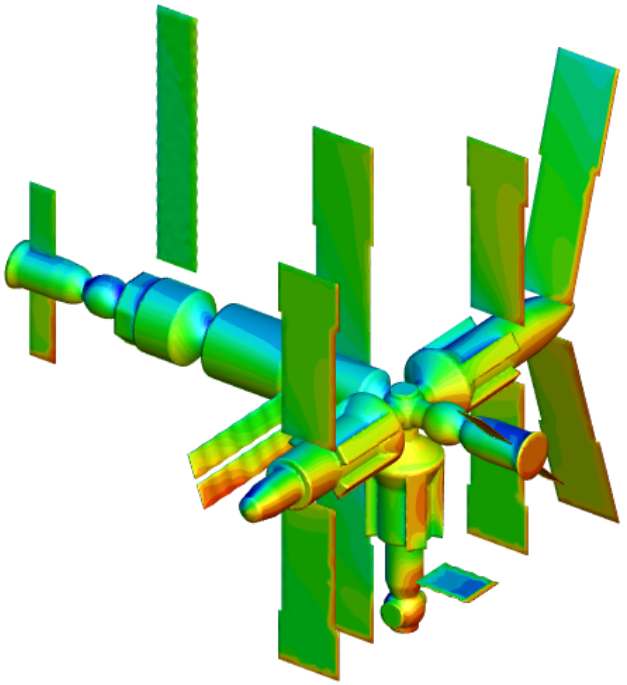

Mir Space Station

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

Simulation of Complicated Shapes

Sandia National Laboratories

P&EM News Note:

A new cut-cell generation algorithm has been developed and deployed in the 3D Direct Simulation Monte Carlo code SPARTA. The algorithm has been demonstrated with a variety of complex 2D shapes, such as the "sea urchin" geometry shown below. SPARTA now computes all the cut cell volumes, identifies any split cells, colors all grid cells as inside, outside, or cut/split. Each surface in a split cell is tagged by which split volume it belongs to, which will be needed for tracking particles into the split cells. Infinitely thin surfaces are detected and correctly dealt with during molecular advection.

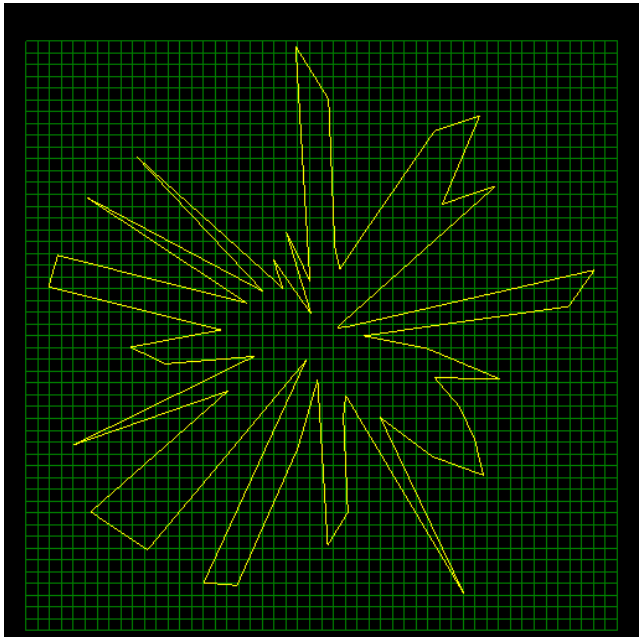

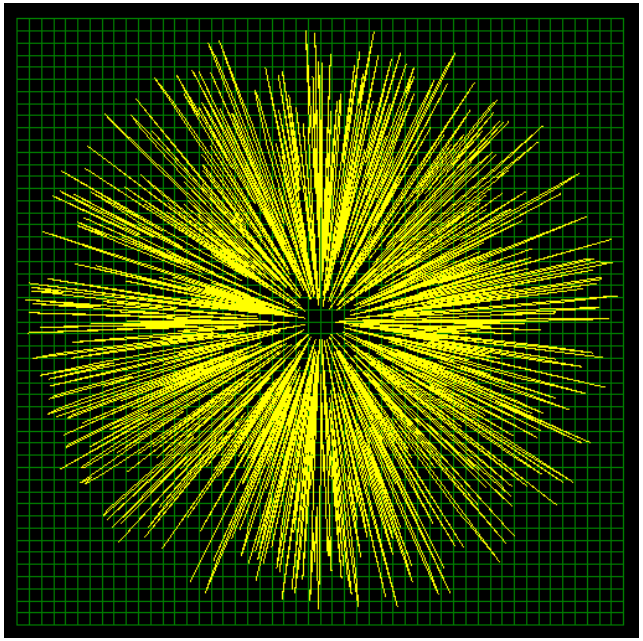

In-Situ Visualization

Sandia National Laboratories

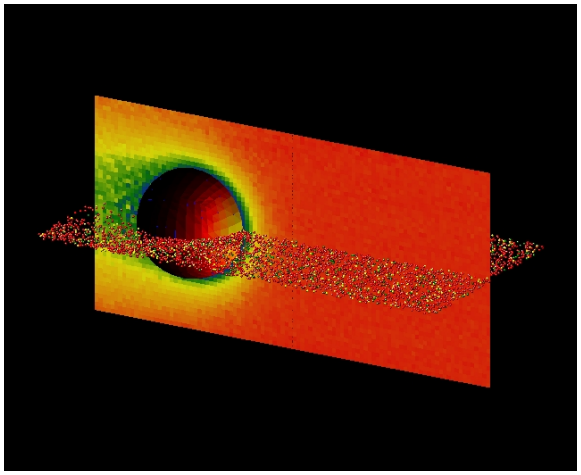

Not a replacement for interactive viz, but ...
Quite useful for debugging & quick analysis
At end of simulation (or during), instant movie

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

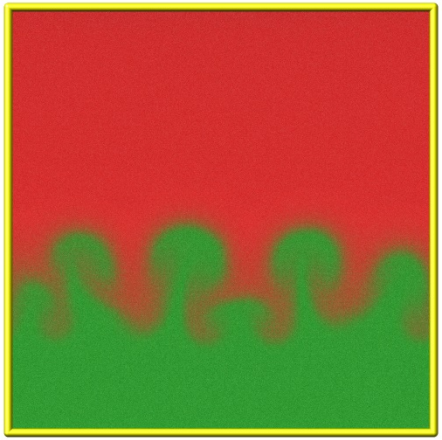

Load Balancing

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

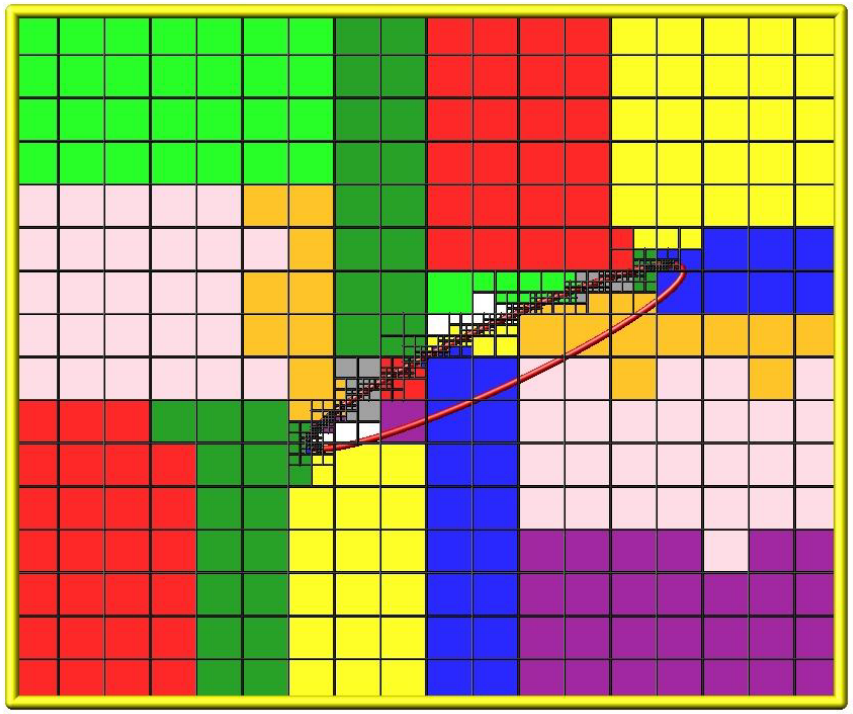

RCB is fast
Bigger cost is data move
Example:
1B cells on1024 BG/Q nodes
Worst case: move all cells
Balance time = 15 s:
(RCB=2, move=12, ghosts=1)

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

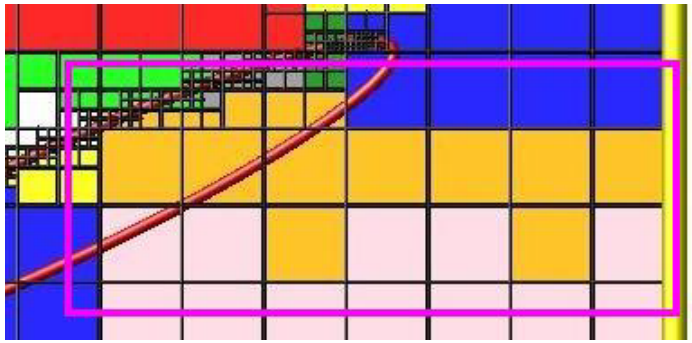

- Efficiently distributes grid information across processors
  - With sufficient cutoff, only one communication per step
  - Multiple passes if needed (or can bound molecule move)
- Communication with modest count of neighbor processors

Work in progress

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

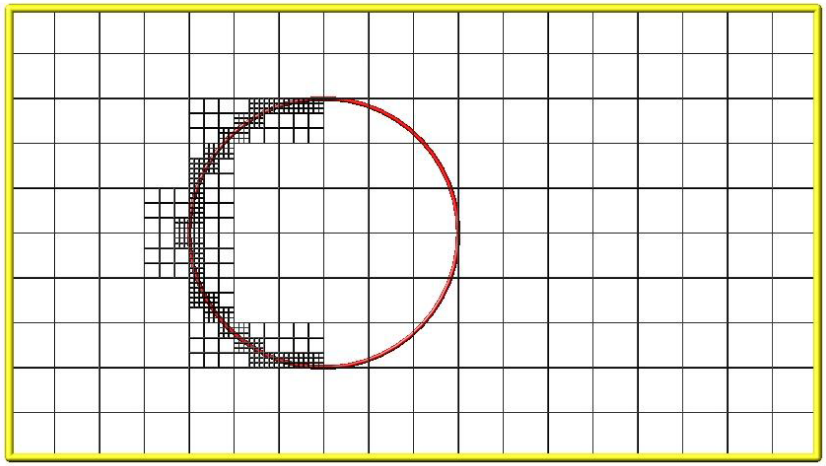

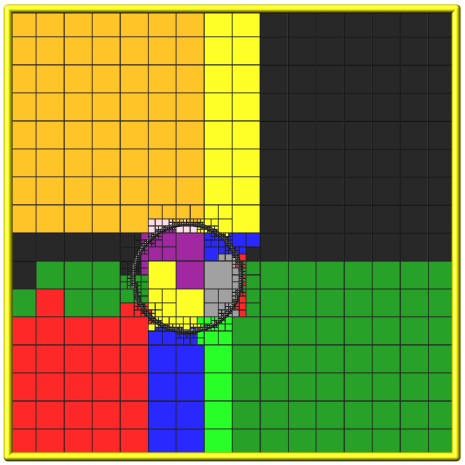

- Another setup task: label cells as outside/inside
- Simple if pre-processing, in situ easier for large problems

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.

Parallel Performance

Sandia National Laboratories

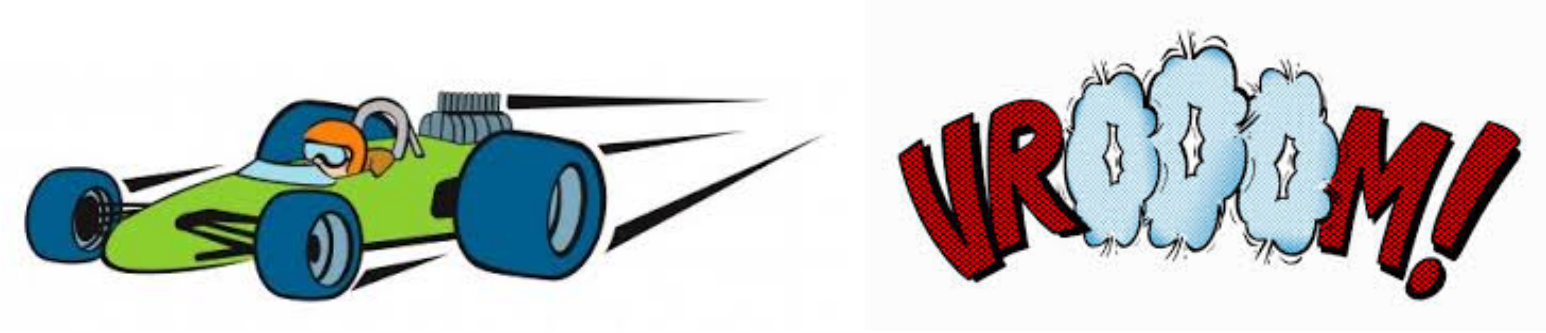

SPARTA Benchmarking

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

Sandia National Laboratories

16 cores/node
1 task/core

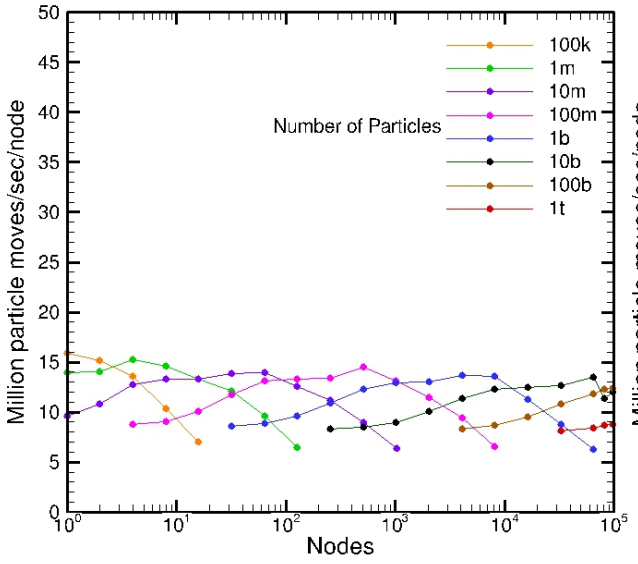

16 cores/node
2 tasks/core

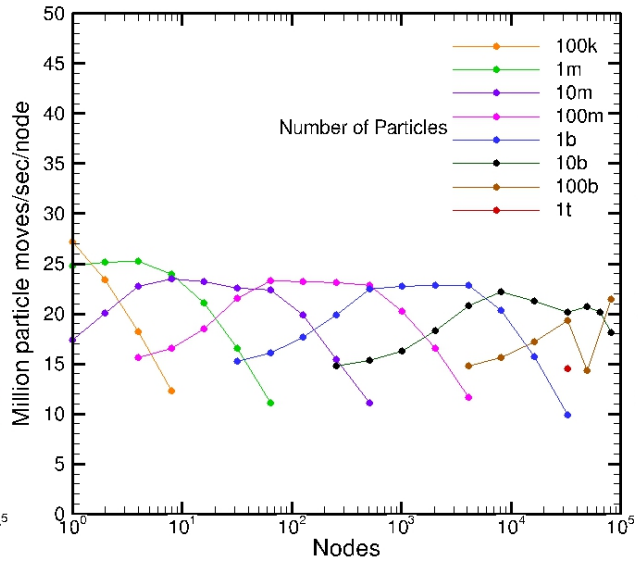

16 cores/node
4 tasks/core

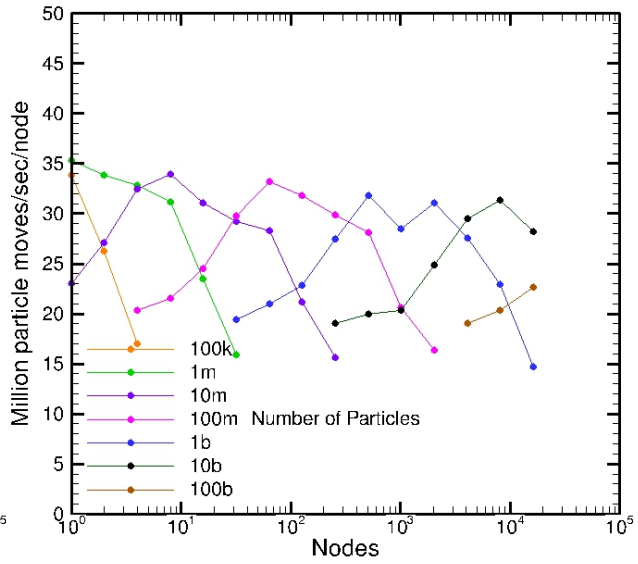

- Weak scaling indicates, 10% peak performance reduction from 1 to 10^6 cores
- 2 tasks/core gives 1.5x speedup, 4 tasks/core gives 2x speedup
- A total of 1 trillion molecules can be simulated on one third of the BG/Q
- Maximum number of tasks is 2.6 million

SPARTA Benchmarking (FM)

16 cores/node, 1 task/core

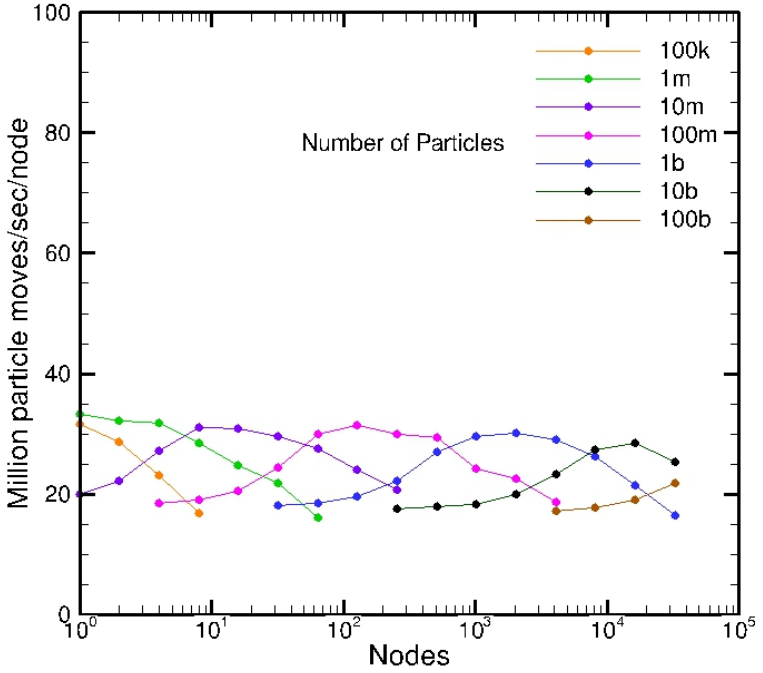

16 cores/node, 4 tasks/core

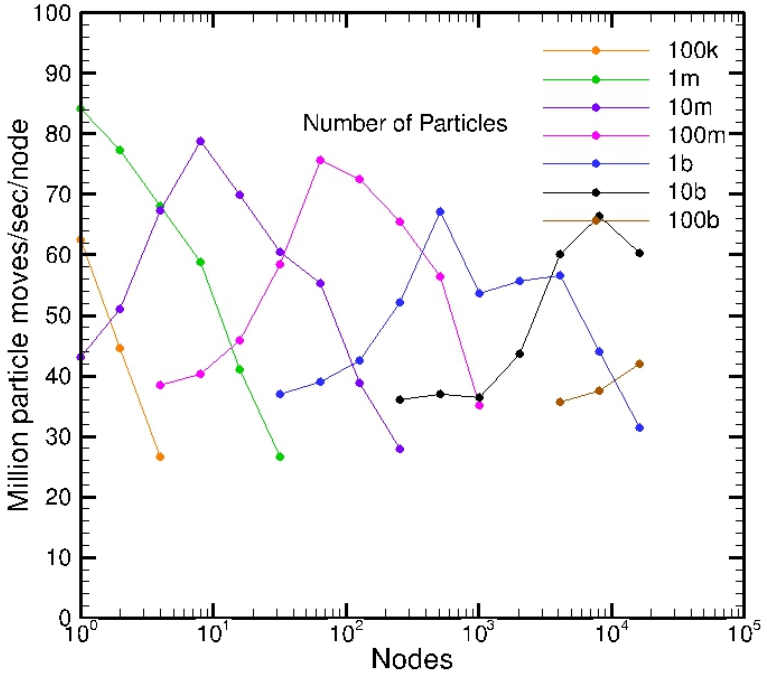

- Free-molecular (FM) calculations stress-test for communications
- 2x speedup compared to collisional

Sparta vs Existing Codes

Sandia National Laboratories

The scalability of the new code framework was compared to the most computationally efficient, existing parallel DSMC codes (3D NASA-DAC, 2D SNL-Icarus). Even for the most severe cases the new code's CPU cost is significantly lower, indicating a more efficient parallel implementation.

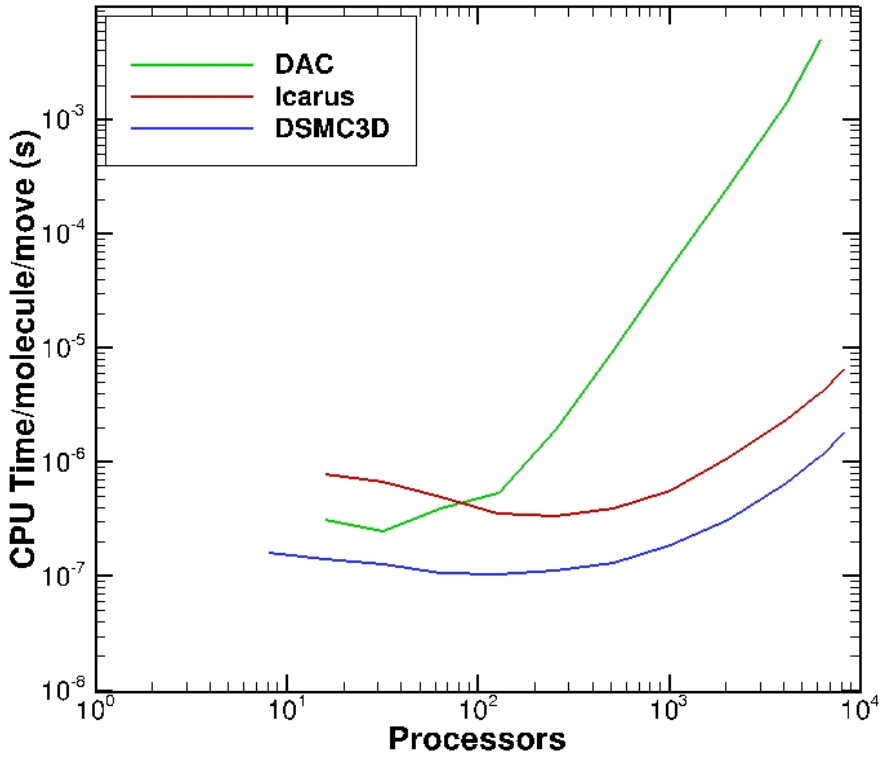

Calculating forces on particles

Sandia National Laboratories

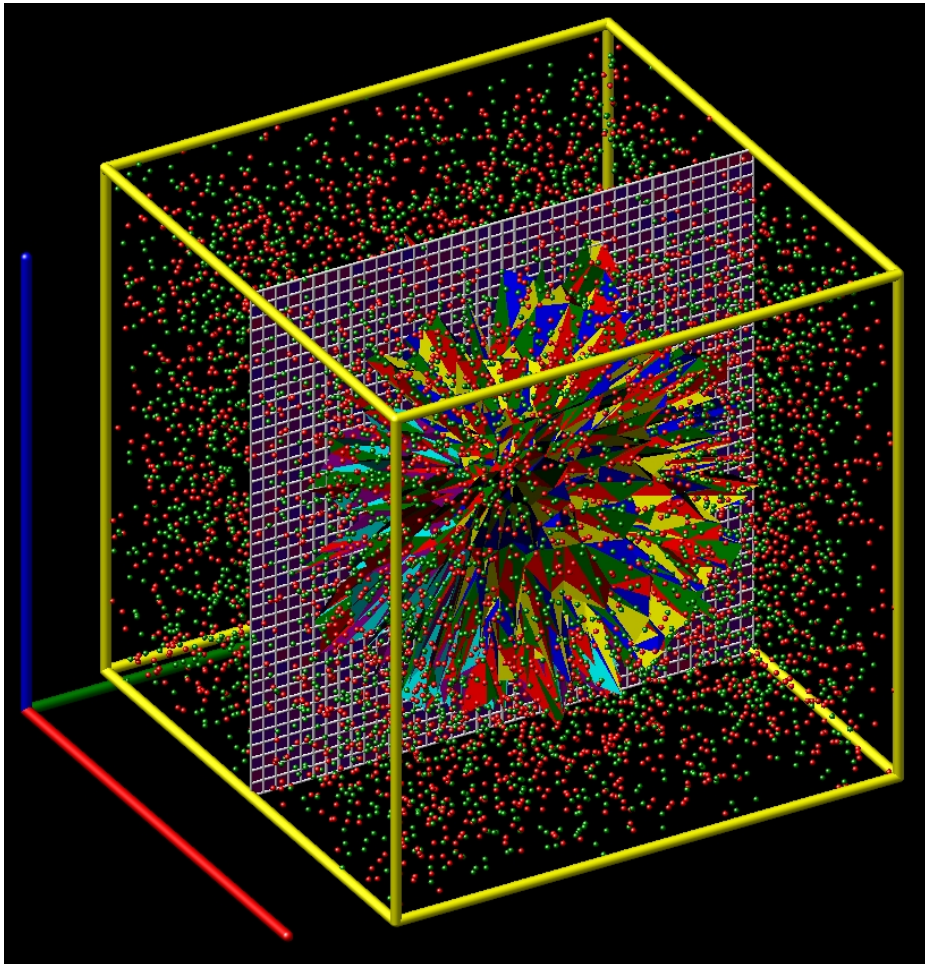

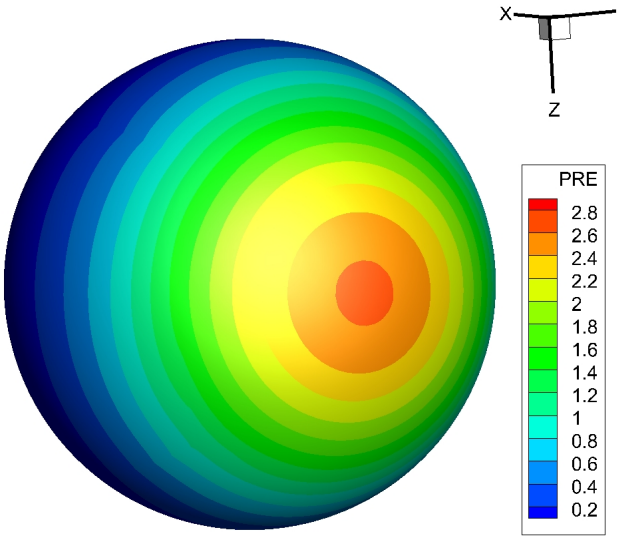

This is SPARTA

Sandia National Laboratories

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

The objective of DSMC is to simulate complicated gas flows using only collision mechanics of simulated molecules in the regime described by the Boltzmann equation

Thermophoretic force between two parallel plates

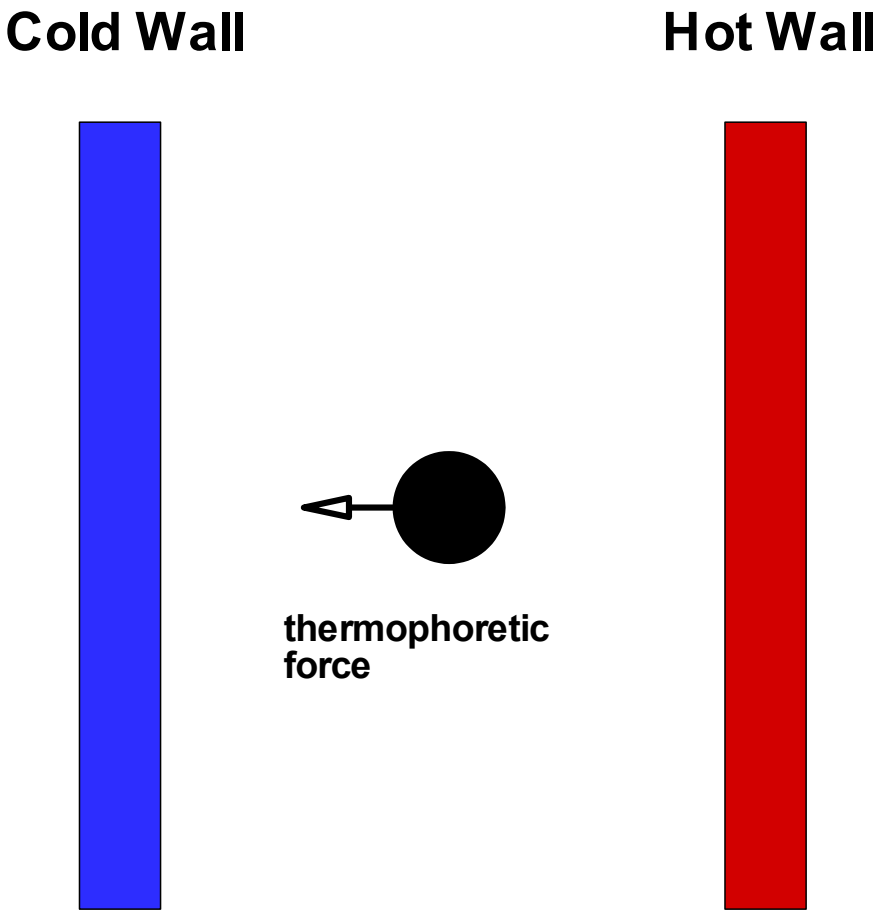

Generation of thermophoresis

Sandia National Laboratories

- Momentum transfer from gas molecules to macroscopic particle.
  - Larger on hot side.
  - Smaller on cold side.
  - Not drag - exists even without net mass flow.
- Heat flux also exists without net mass flow.
- Thermophoresis is associated with the heat flux.

Test case conditions

Sandia National Laboratories

Gas	Argon, Helium
Pressure	0.1–1000 mTorr
Temperature	263-283 K
Domain length	0.01 m
Grid	100 x 10 (Low densities) 1000 x 10 (High densities)
Time step	10^{-7} (Low densities) 10^{-8} (High densities)
Molecules per cell	100
Molecular model	VSS

Thermophoretic force as a function of pressure

Argon

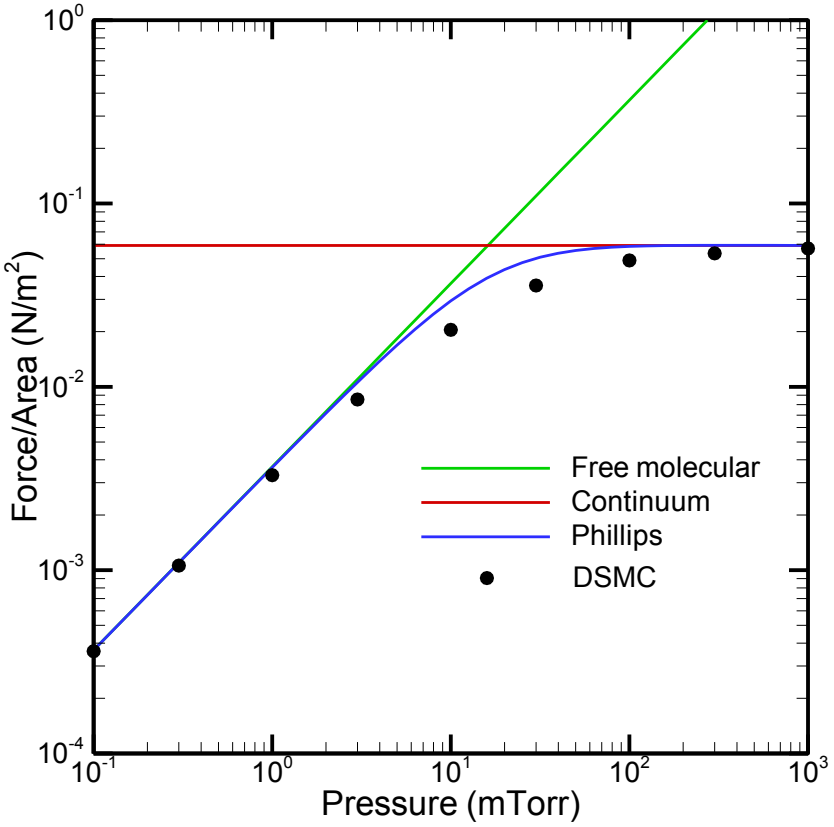

Helium

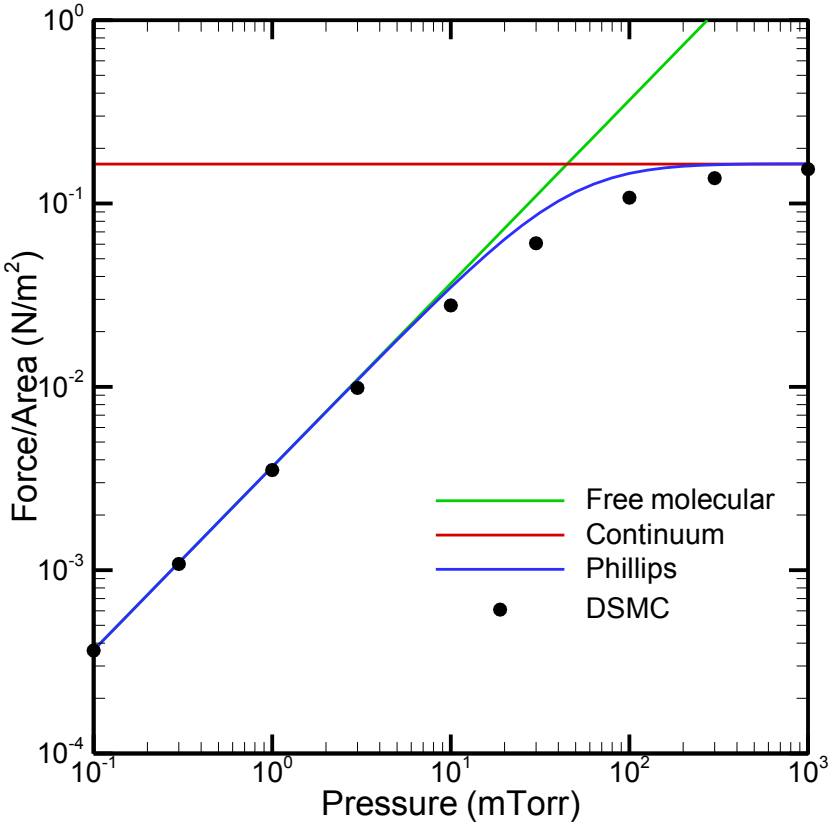

Microscopic Dissociation Model

Sandia National Laboratories

When the vibrational level of a molecule exceeds the dissociation limit, dissociation occurs.

Maximum reachable vibrational state: $i_{\max} = \text{int}[E_c / (k\Theta_v)]$

where: $E_c = E_{trans,pair} + E_{vib,mol}$

Dissociation occurs when: $i_{\max} > i_d$

Dissociation rate coefficient:

$$k(T_{tr}, T_{vib}) = \frac{2\sigma_{ref}}{\varepsilon\sqrt{\pi}} \left(\frac{T_{tr}}{T_{ref}} \right)^{1-\omega} \left(\frac{2k_B T_{ref}}{m_r} \right)^{1/2} \left\{ \sum_{i=0}^{i_d} Q\left[\frac{5}{2} - \omega, \frac{\Theta_d - (i-1)\Theta_v}{T_{tr}} \right] \frac{\exp[-i\Theta_v / T_{vib}]}{z_{vib}(T_{vib})} + B \sum_{i=i_{d+1}}^{\infty} \frac{\exp[-i\Theta_v / T_{vib}]}{z_{vib}(T_{vib})} \right\}$$

Comparison to Experimental Data

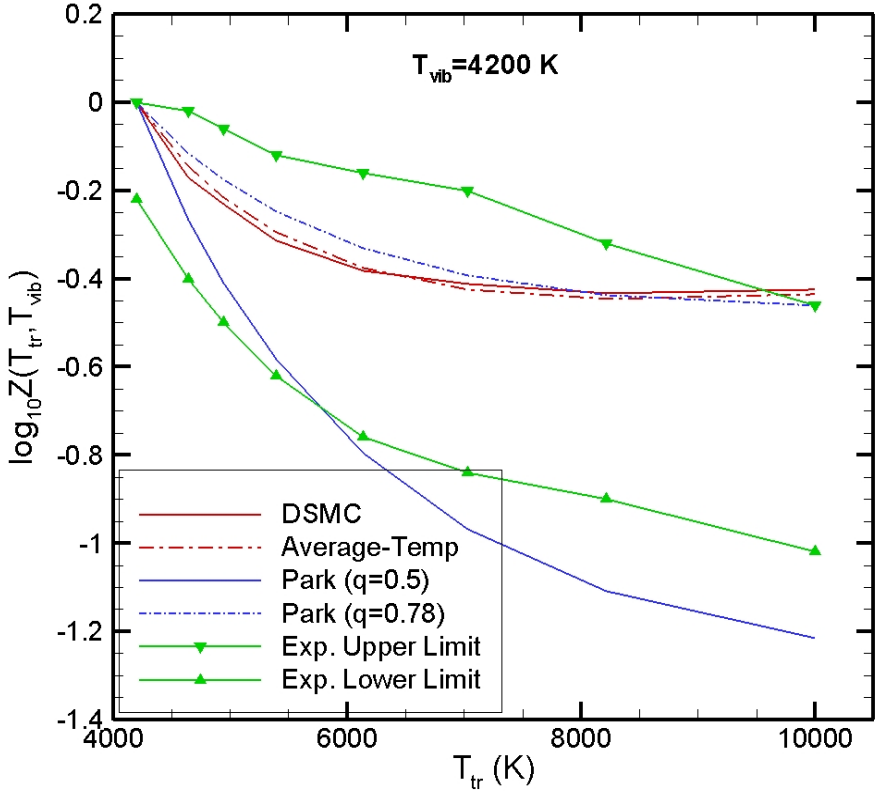

In the experiment (Sergievskaya et al.), T_{vib} was constant at 4200 K. DSMC is in agreement with measured non-equilibrium reaction rates.

Park's model with q=0.78 (for low T) is in good agreement with data.