SAND2010-0966C

A tale of two scales:

Petascale computing meets multiscale simulations

Helgi Adalsteinsson
Scalable Computing R&D Dept.
Sandia National Laboratories
Livermore, CA

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy's National Nuclear Security Administration
under contract DE-AC04-94AL85000.

Overview

- Motivation
  - How will you use millions of threads?
  - Redefining the meaning of “scalable”
  - High performance computing and the quest for better science
- Multiscale multiphysics simulations
  - Dr. Frankenstein’s simulation system
  - Better software engineering (for given definitions of “better”, “software”, and “engineering”).
  - What’s going wrong now?
- The quest for improved fidelity
  - You too can contribute to global warming
- Conclusions

HPC systems are growing…

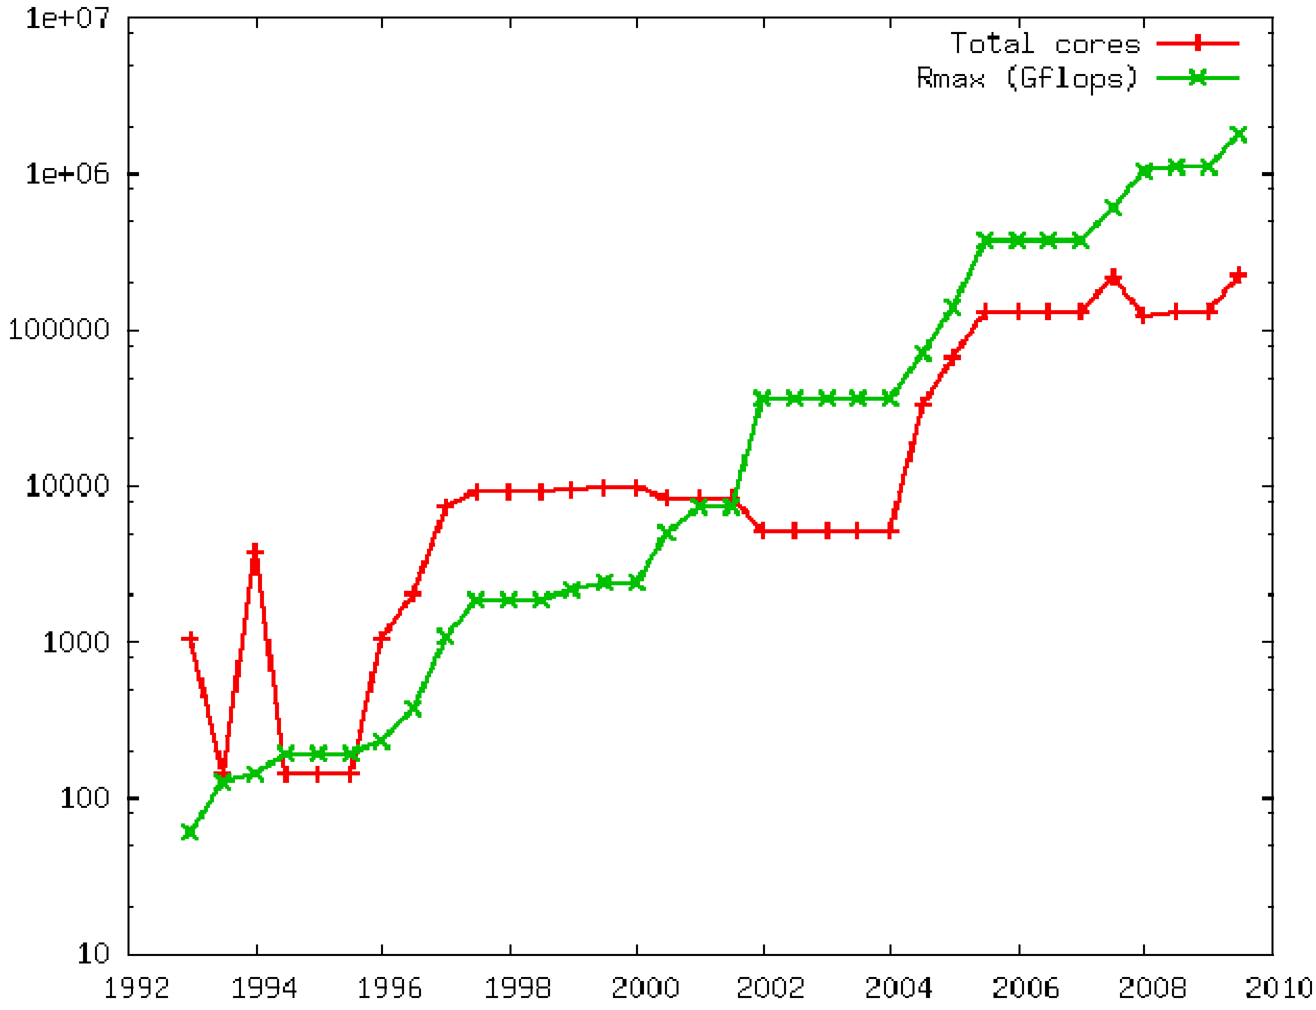

…and HPC applications can't keep up

- Very few "hero run" simulations can utilize the full scale of today's capability systems
- (Almost) no science apps run "at scale" on current capability systems

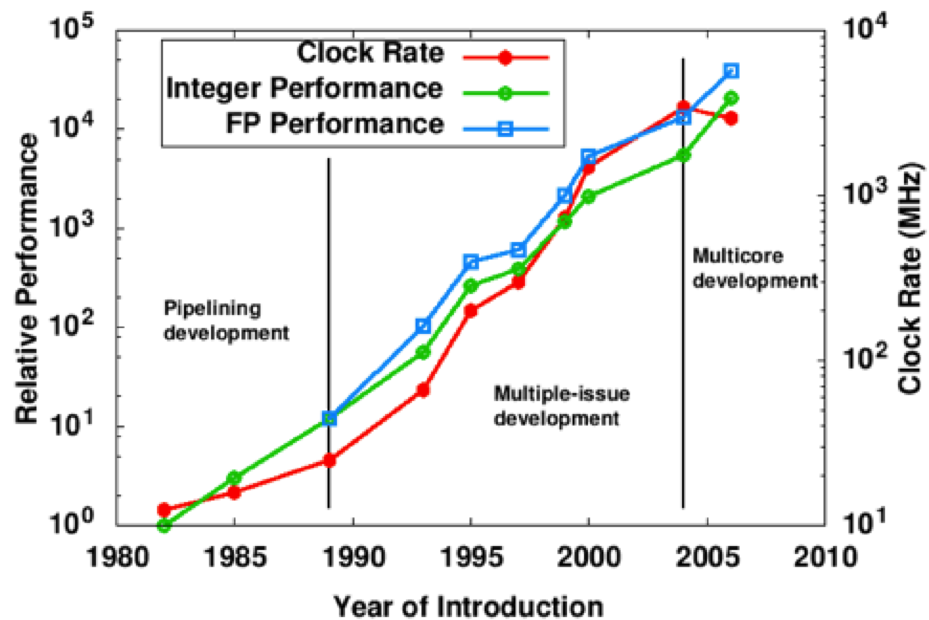

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Where does scalability come from?

- We can envision our system as a stack
  - Application (and associated libraries)
  - Middeware (runtime systems and programming models)
  - Hardware, firmware, and OS
- The stages in the stack are not clearly delineated
  - Our group is heavily involved in the middleware and library layers
  - Tomorrow's middleware needs to exist today

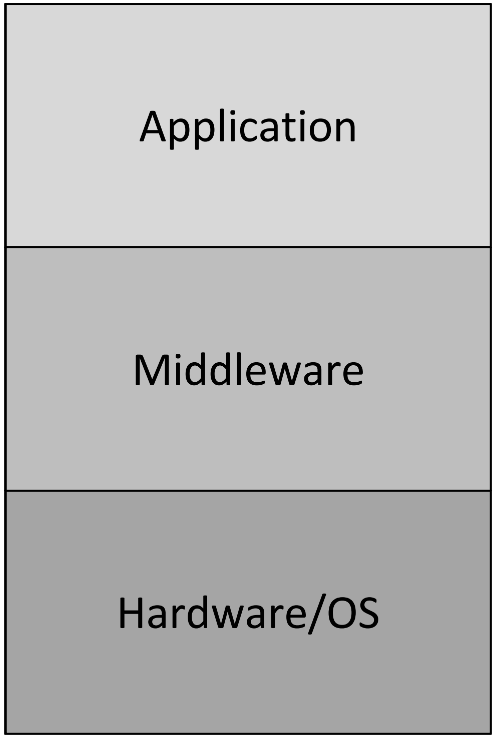

A step back: What do we mean by "scale"

- Traditional scalability definitions are biased toward hardware/OS improvements
  - Strong scaling: Fixed total problem size
  - Weak scaling: Fixed per-processor problem size
- The more significant challenge is what we expect to get from the application
  - How can we get more valuable results?
  - What can we do to enable new classes of applications?

The case for multiscale multiphysics simulations

- You find yourself with a system that:
  - has regions that are characterized by disparate-scale physics
  - … and those system scales are bi-directionally coupled
  - … but yet separated in a way that allows the system to be decomposed onto two or more simulation models
- … so you decide to use more than one simulation model together to get fine-scale accuracy for a coarse-scale price
  - (your idea is not new – concepts dating back to 1896)
- … and now all that remains is stitching the thing together

Component design for multiscale

- (or design patterns, if you insist)
- Set out to identify and design software components for atomistic-to-continuum multiscale
  - Is there room for a shared middleware layer?
  - Can we identify sensible component interfaces for information sharing?
  - Can we make multiscale development more sensible?

First case study: Multiscale electroosmotic flow

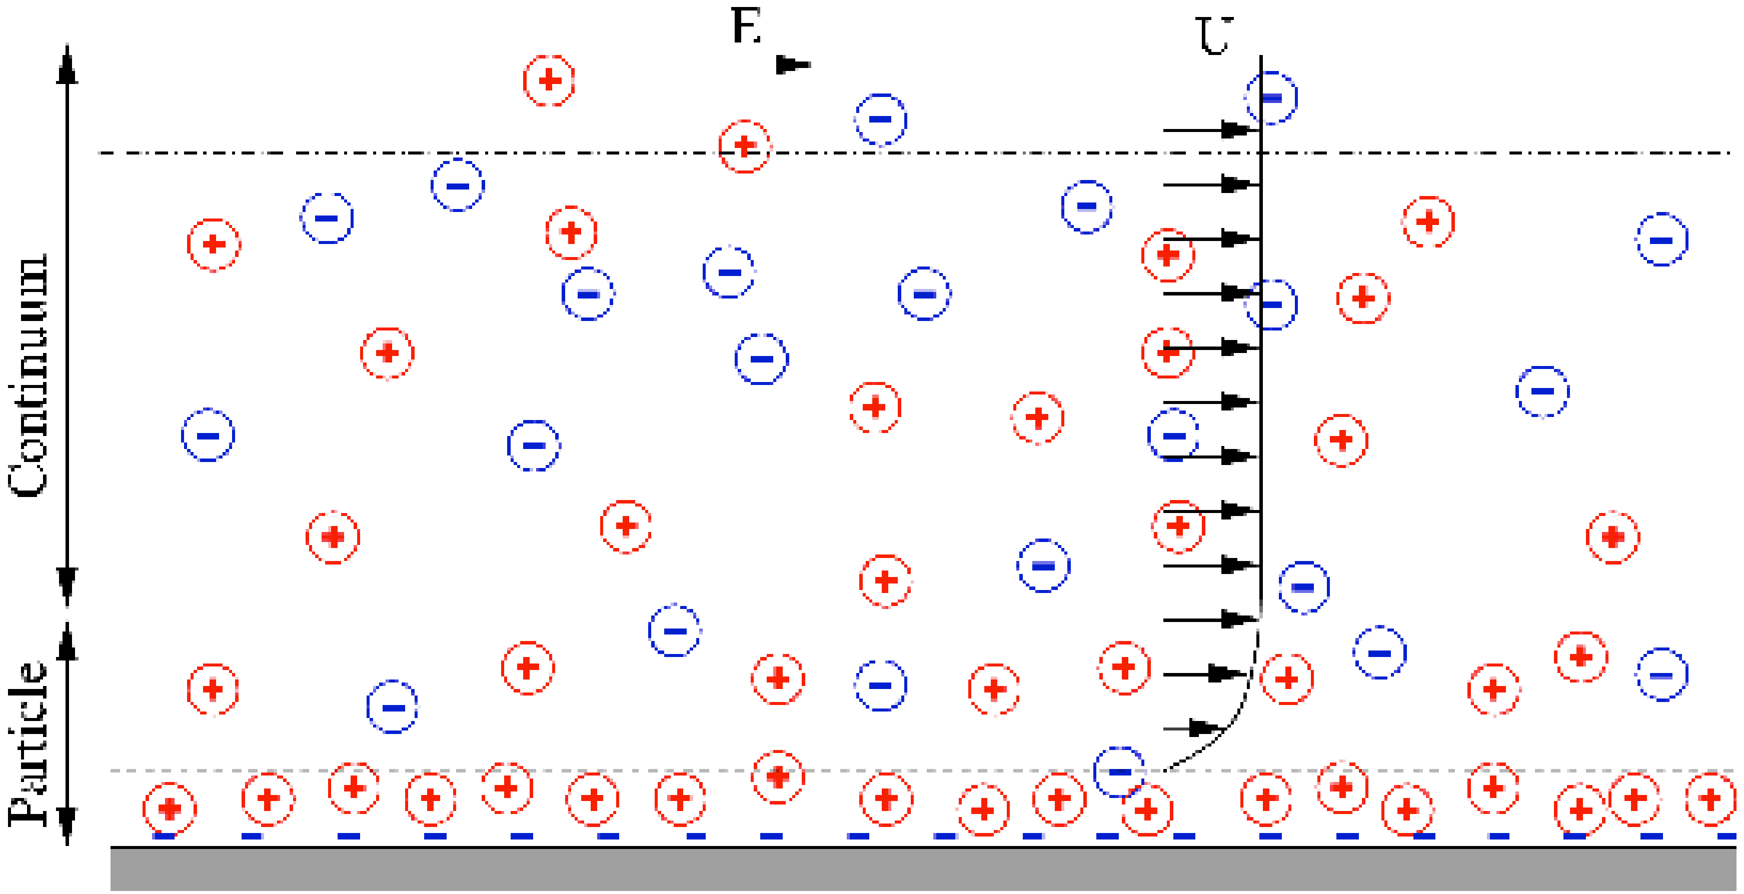

The power of small

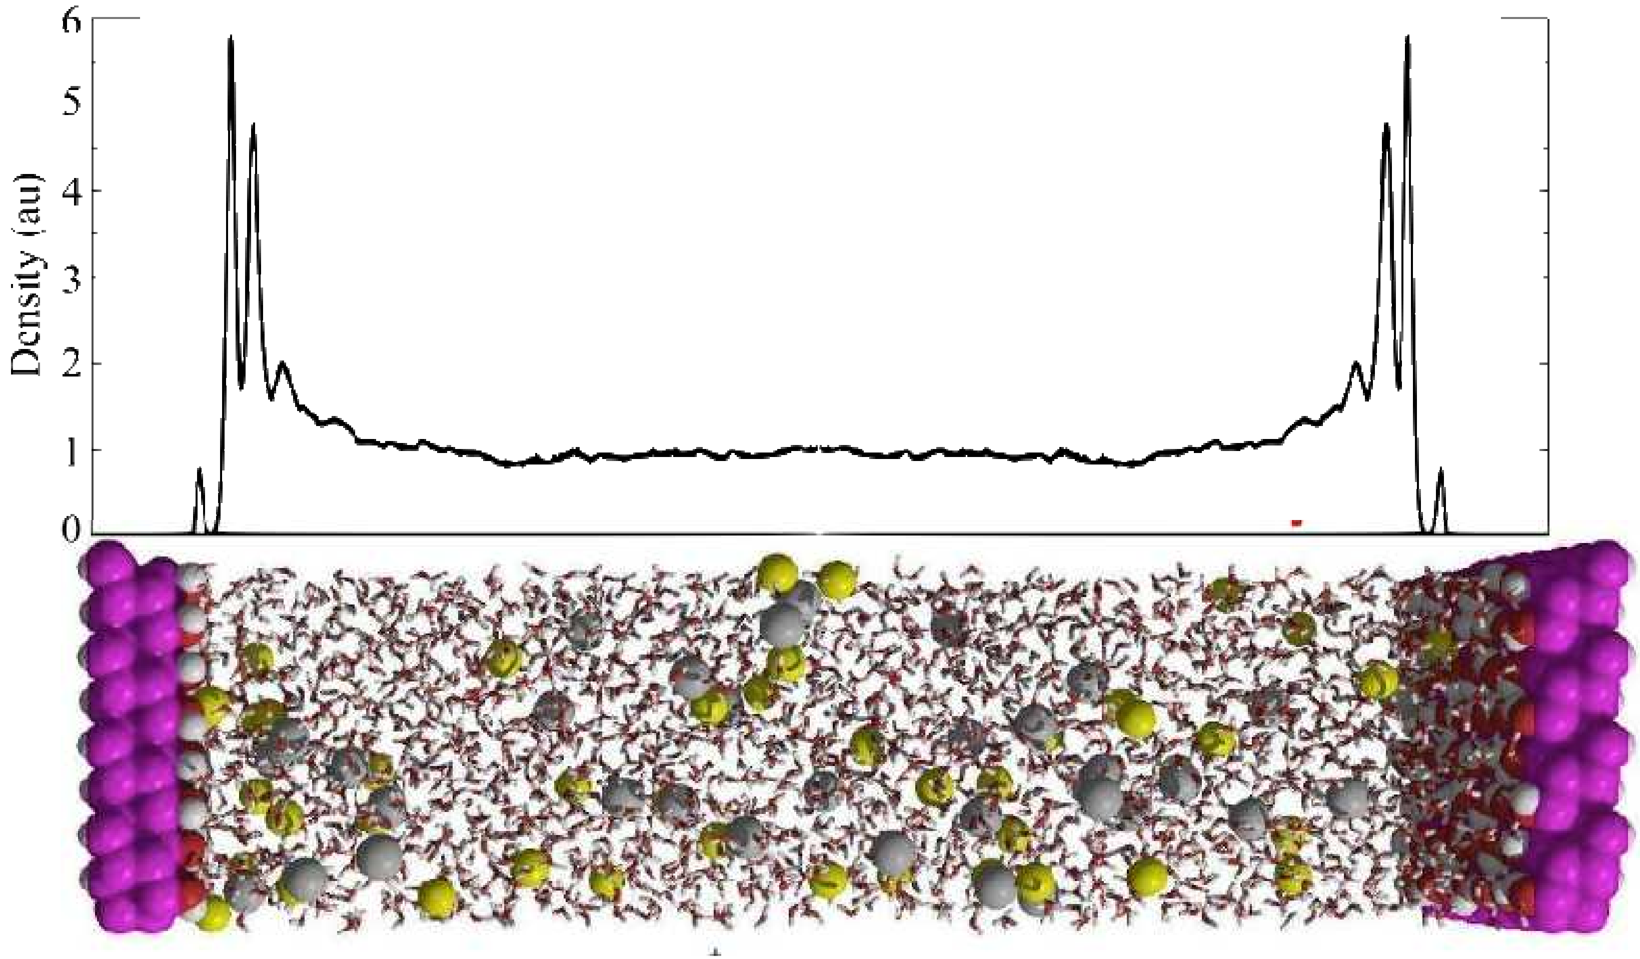

Look who's driving

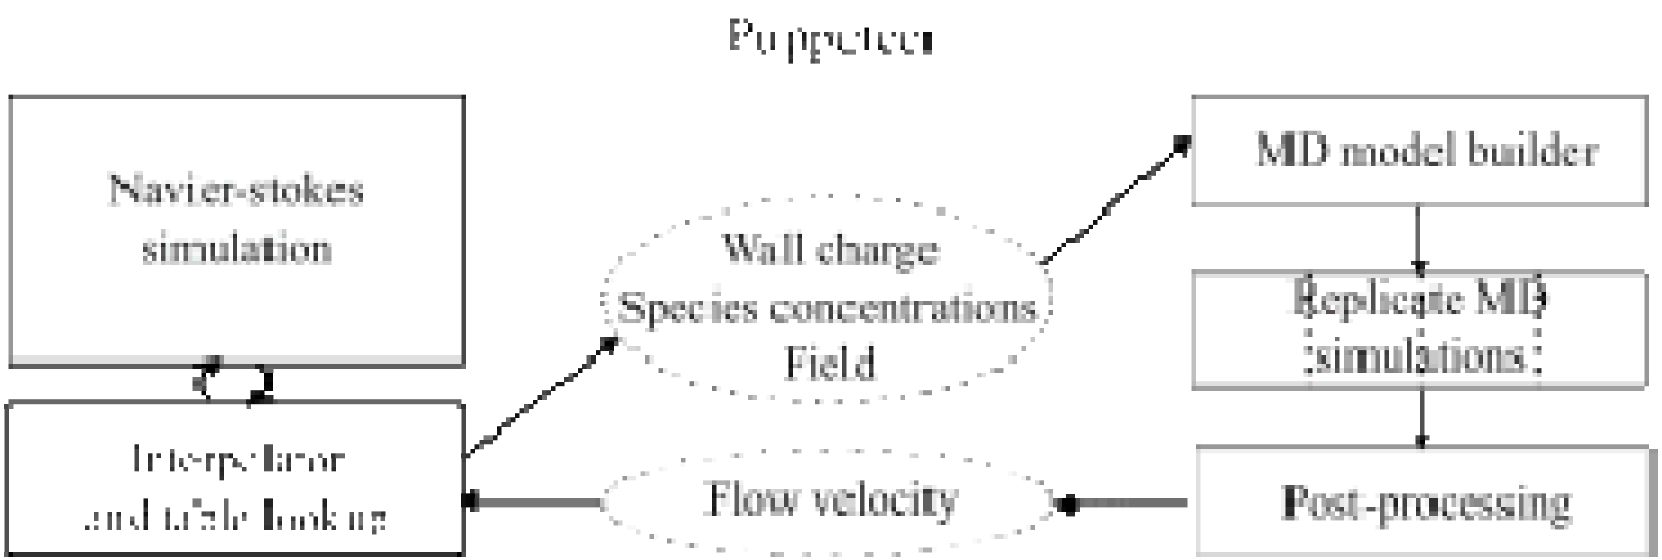

A simplified example

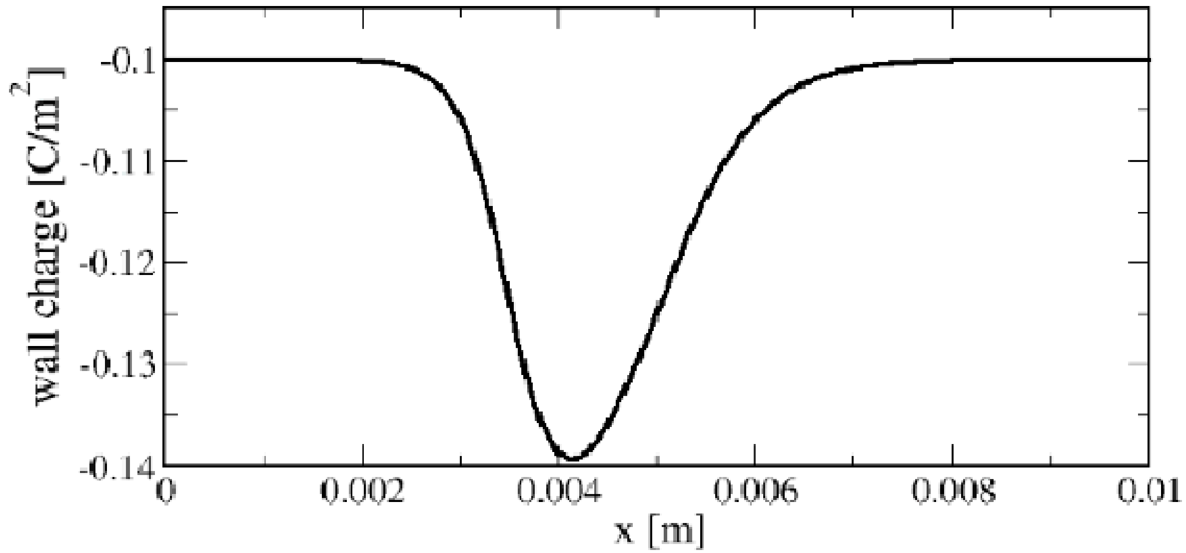

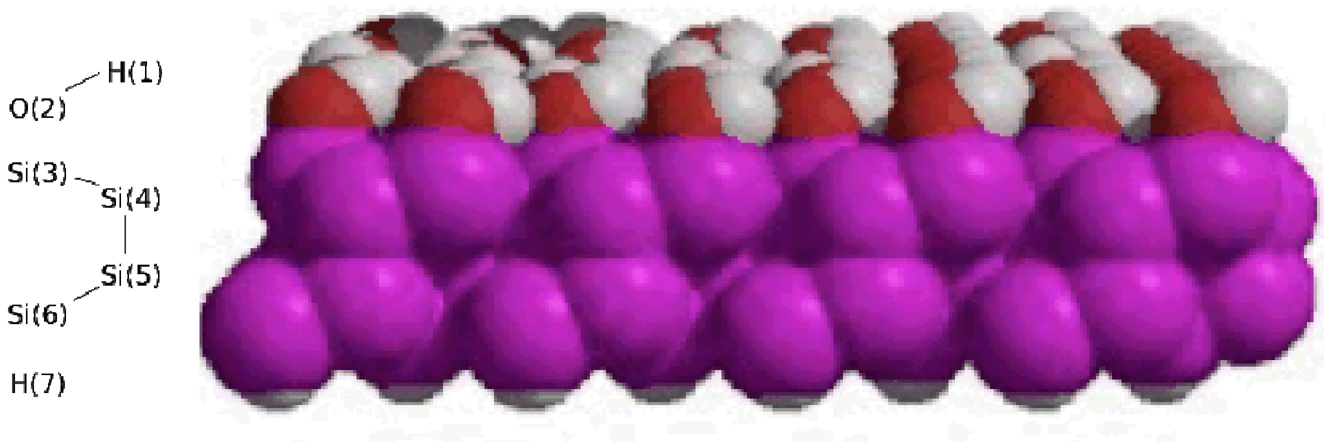

And beautiful results emerge*

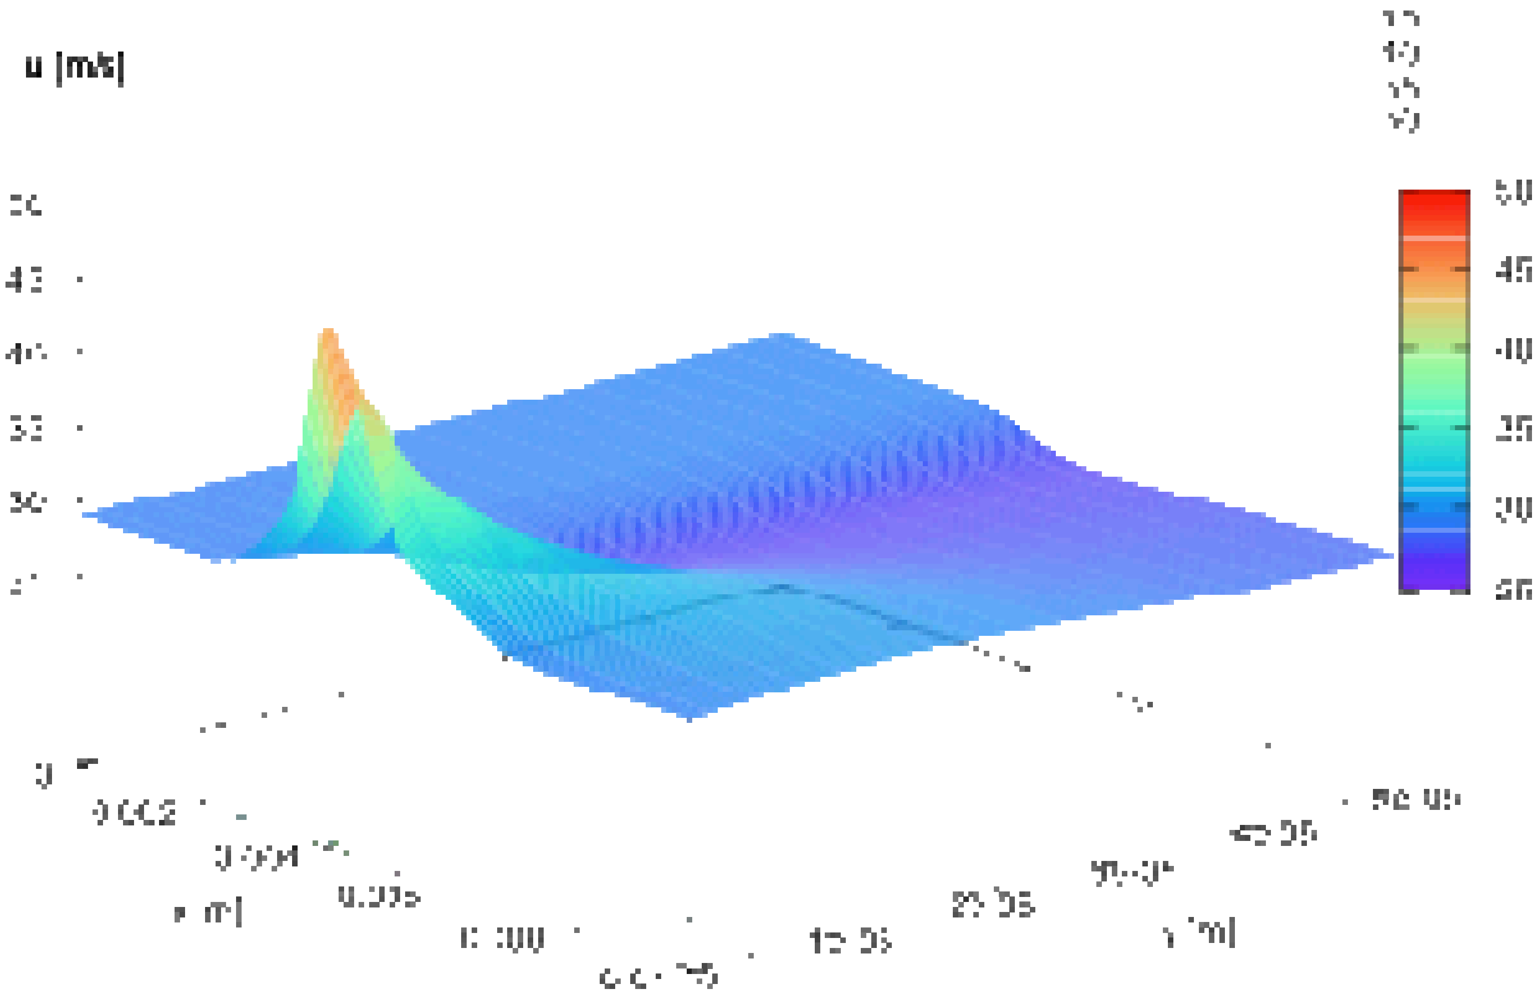

Second case study: Multiscale electrophoresis

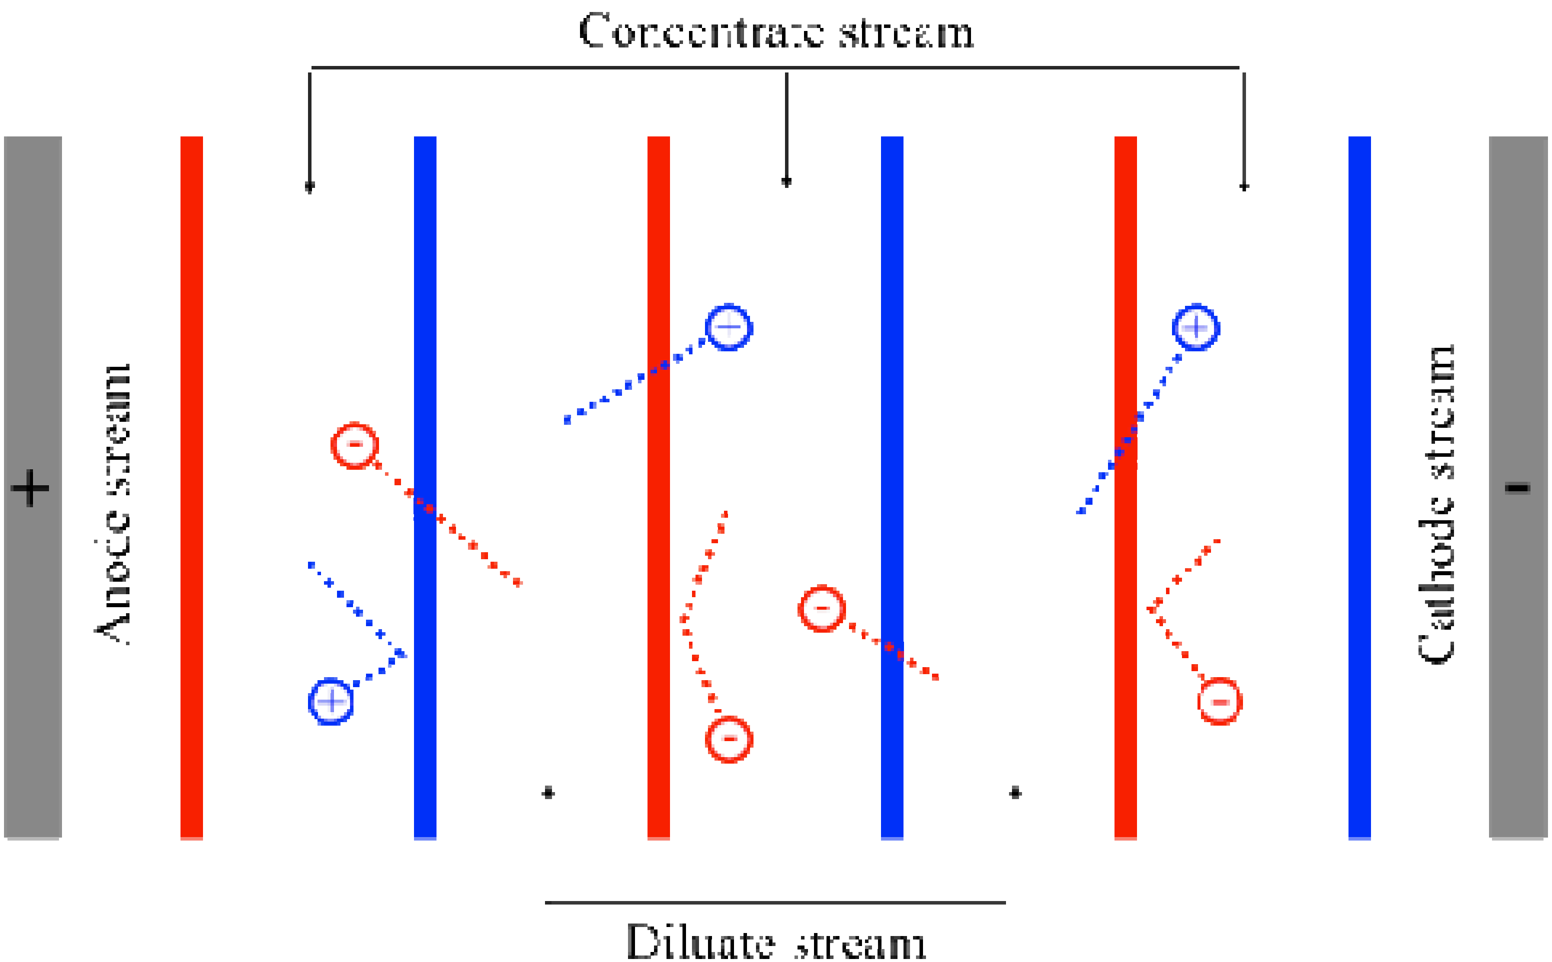

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy's National Nuclear Security Administration
under contract DE-AC04-94AL85000.

The basic simulation system

Look who's driving now

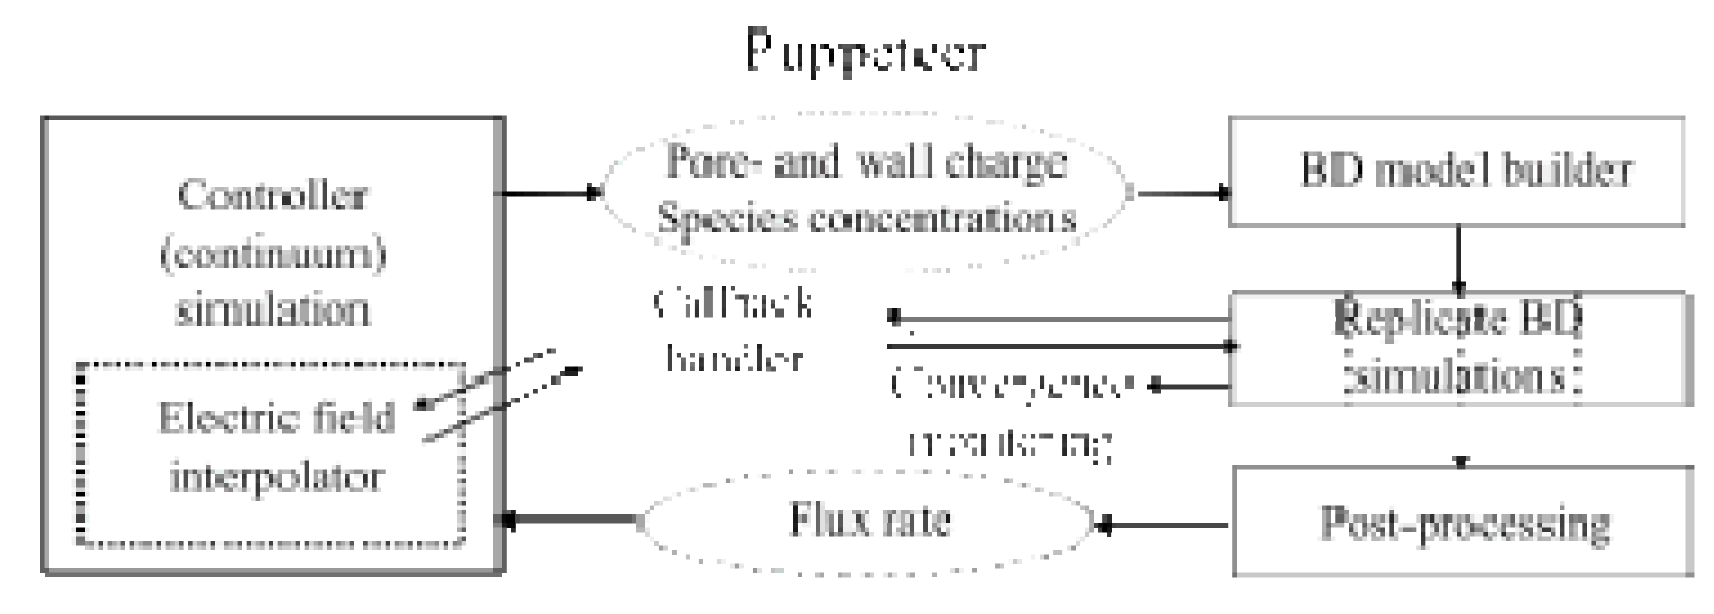

Parallelization

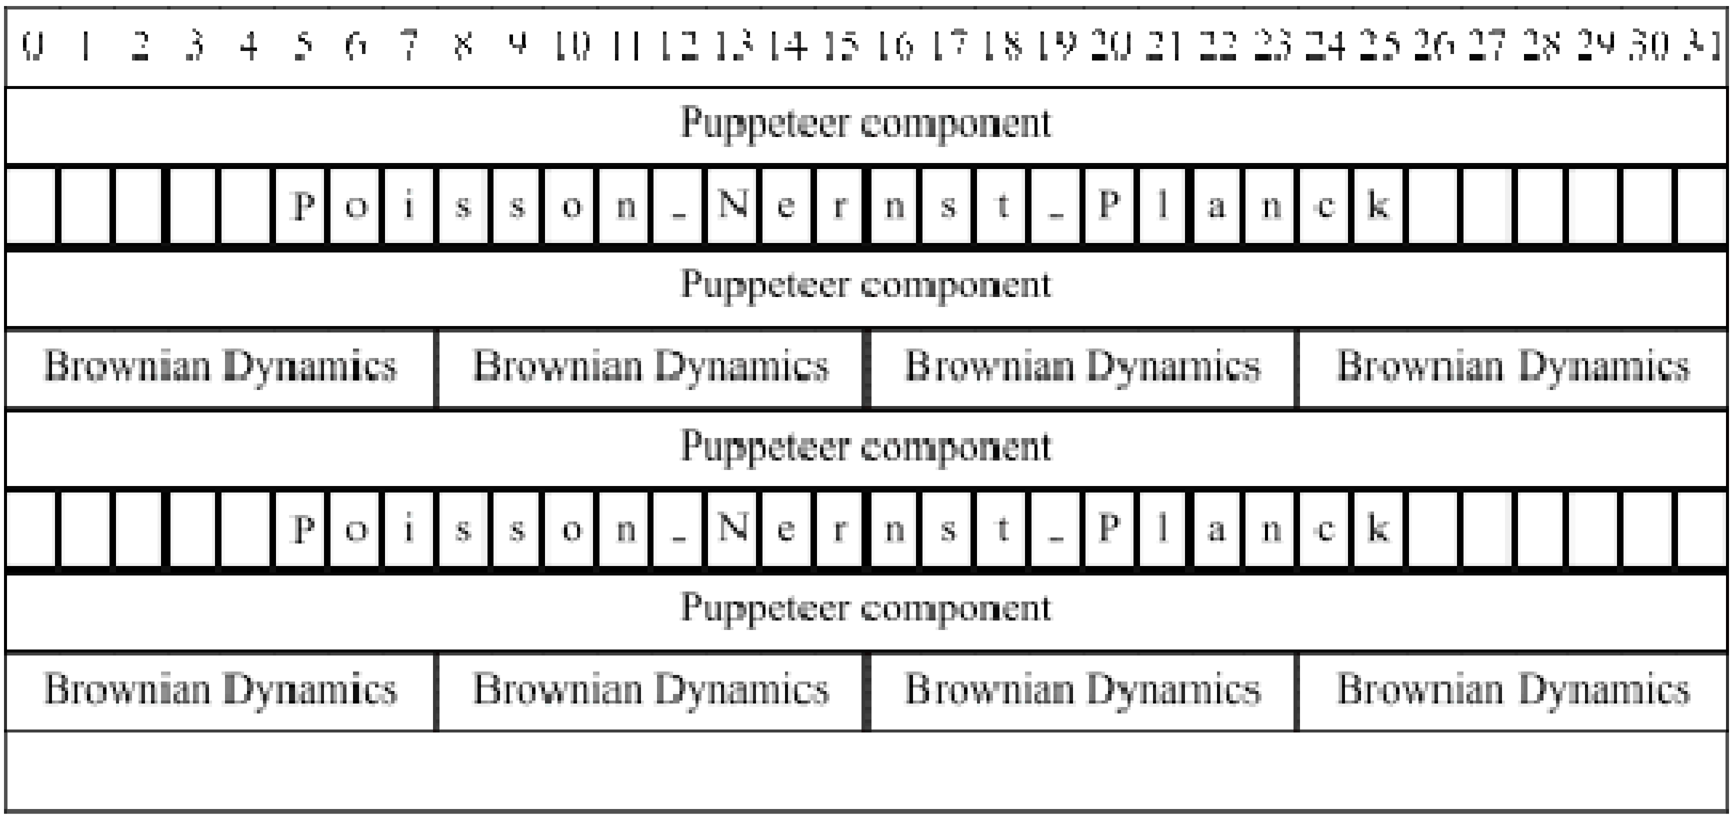

And all is well*

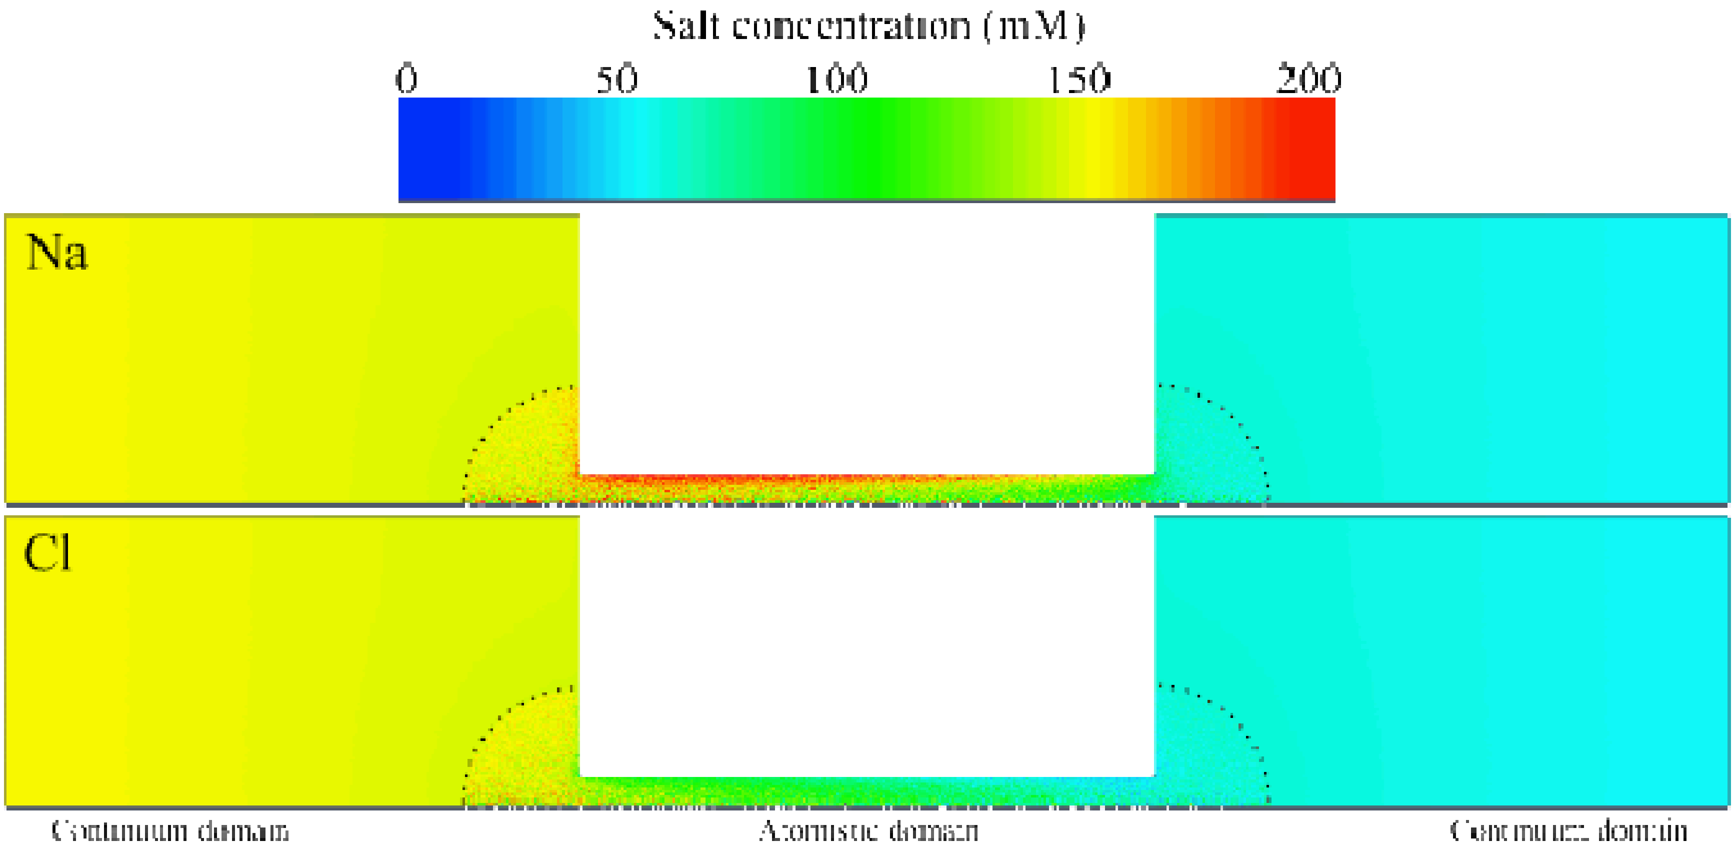

*Living in the under-resolved lane

- Each particle simulation is under-resolved
  - … but they may not be identically under-resolved
  - … and we have no mechanism to assess the quality of the final simulation

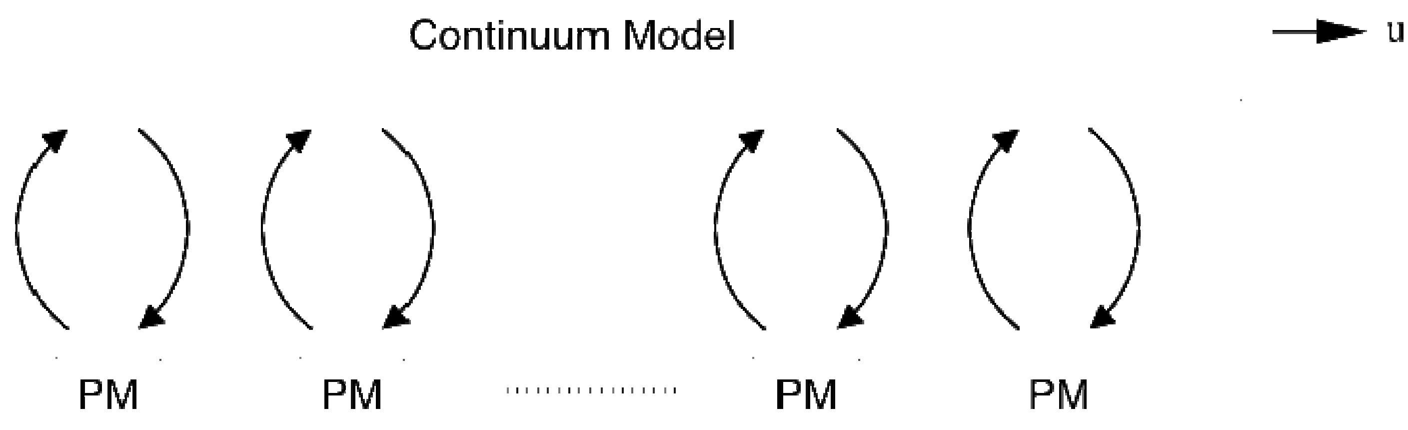

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Uncertainty quantification for the multiscale world

- We need a mechanism to:
  - Assess overall simulation quality
  - Achieve balance between the different parts of a simulation

After all, how hard can it be?

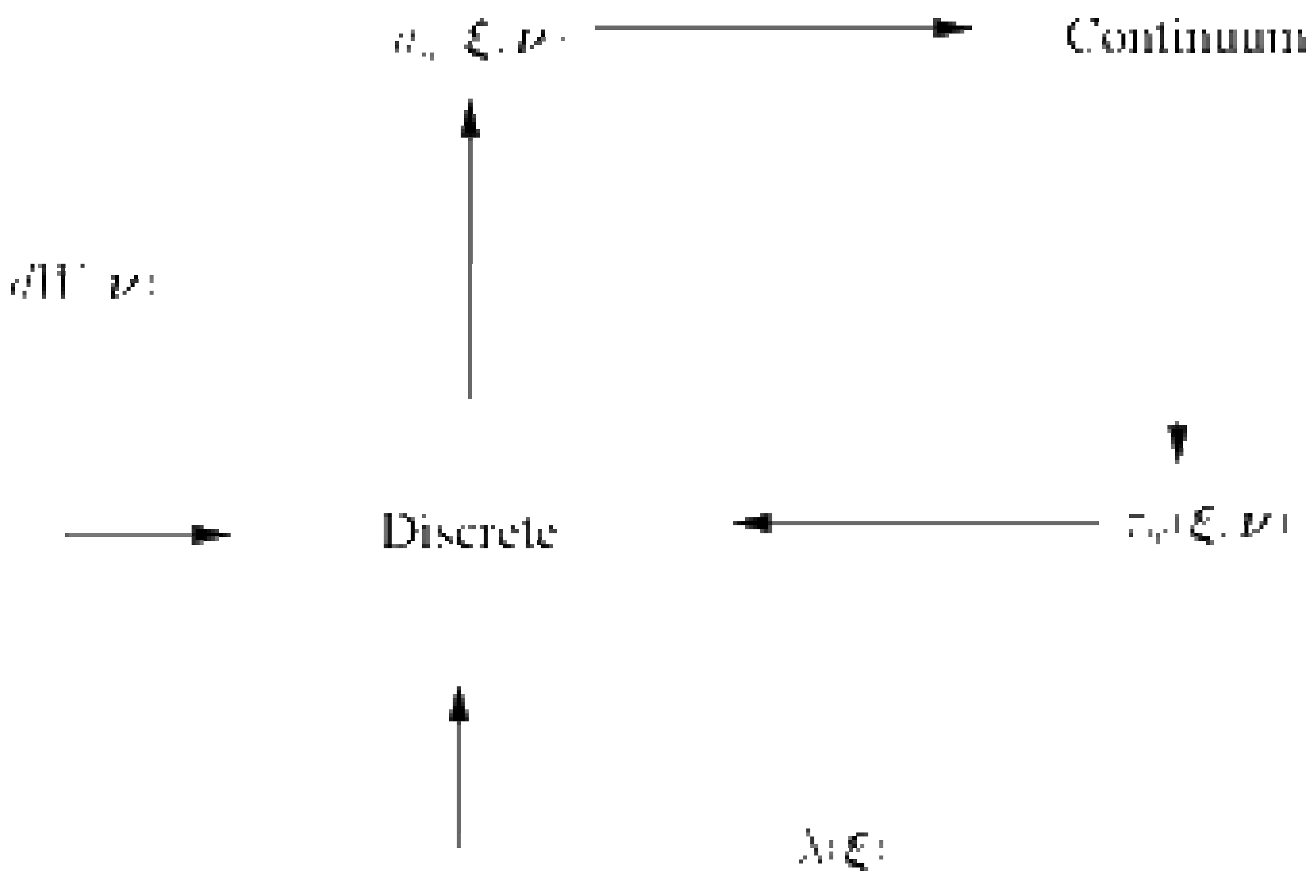

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy's National Nuclear Security Administration
under contract DE-AC04-94AL85000.

…and immediately, there is trouble

- MD systems are a terrible mess
  - Little or no error/uncertainty information available for model parameters
  - Many additional simulation parameters with nonlinear (or completely counterintuitive) effects
  - Noisy stochastic simulations
  - Slow convergence and long simulation times
- Even water is hard
  - Dozens of water models in use…
  - …none of which gets all bulk properties right

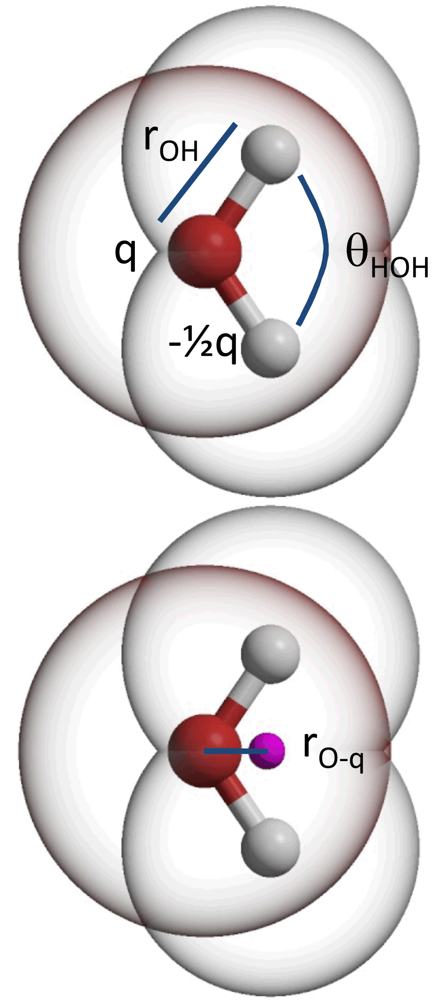

NNSA National Nuclear Security Administration

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

And if you thought water was hard…

If a simple box of water will not yield a nicely behaved simulation…

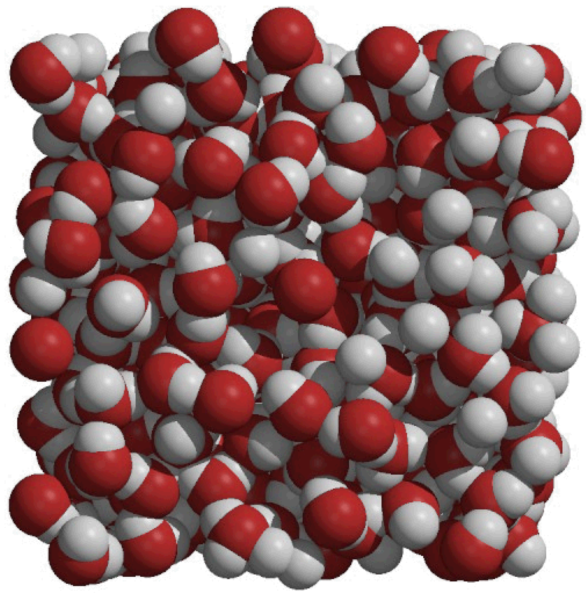

…what about systems that may be worthy of a multiscale simulation

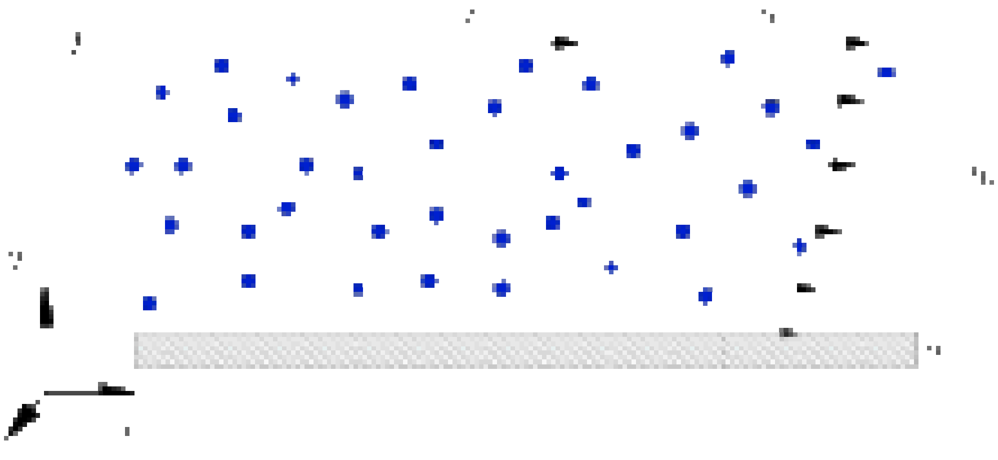

Current status

- Working on pure MD approaches
- Working toward a coupled multiscale Couette simulation

Acknowledgments

- Bert Debusschere (SNL)
- Habib Najm (SNL)
- Reese Jones (SNL)
- Robert Berry (SNL)
- Prof. Kevin Long (TTU)
- Prof. Omar Knio (JHU)
- Maher Salloum (JHU)
- Francesco Rizzi (JHU)
- Prof. Youssef Marzouk (MIT)
- Prof. Michael Hickner (PSU)
- Frank Zendejas (SNL)