GPU TECHNOLOGY CONFERENCE

Klaus Schulten

Department of Physics and Theoretical and Computational Biophysics Group
University of Illinois at Urbana-Champaign

GTC, San Jose Convention Center, CA | Sept. 20–23, 2010

PRESENTED BY NVIDIA.

GPU and the Computational Microscope

Accuracy • Speed-up • Unprecedented Scale

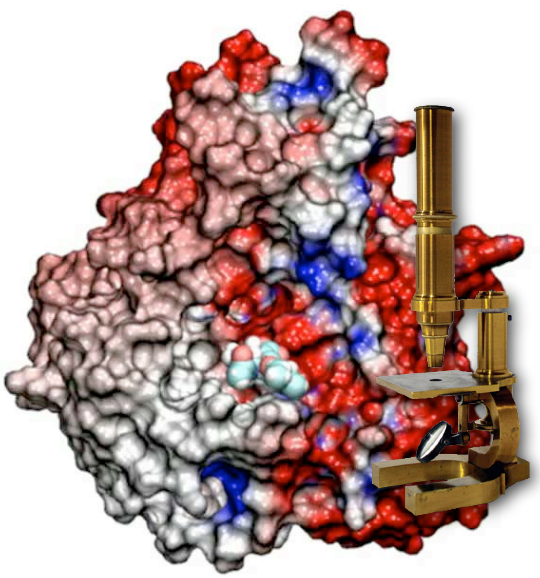

Investigation of drug (Tamiflu) resistance of the "swine" flu virus demanded **fast response!**

Computational Microscope Views at Atomic Resolution...

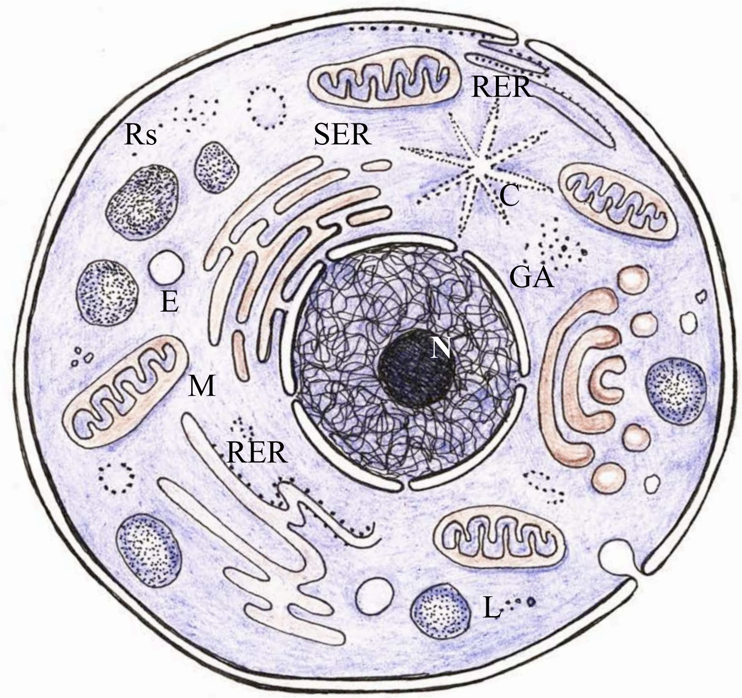

...how living cells maintain health and battle disease

Our Microscope is Made of...

Chemistry

$$U(\vec{R}) = \underbrace{\sum_{bonds} k_i^{bond}(r_i - r_0)^2}_{U_{bond}} + \underbrace{\sum_{angles} k_i^{angle}(\theta_i - \theta_0)^2}_{U_{angle}} + \underbrace{\sum_{dihedrals} k_i^{dihe}[1 + \cos(n_i\phi_i + \delta_i)]}_{U_{dihedral}} + \underbrace{\sum_i \sum_{j \neq i} 4\epsilon_{ij}\left[\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right] + \sum_i \sum_{j \neq i} \frac{q_i q_j}{\epsilon r_{ij}}}_{U_{nonbond}}$$

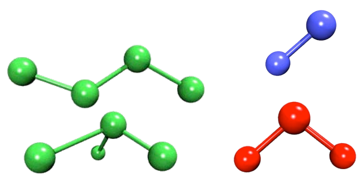

NAMD Software

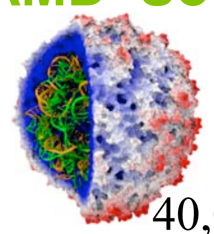

Virus

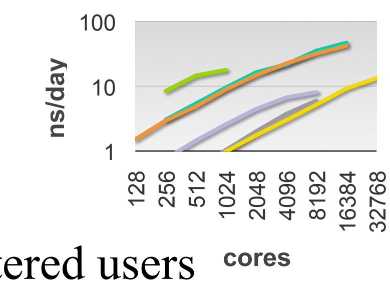

40,000 registered users

Physics

$$m_i \frac{d^2 \vec{r}_i}{dt^2} = \vec{F}_i = -\vec{\nabla} U(\vec{R})$$

Math

$$\vec{r}_i(t + \Delta t) = 2\vec{r}_i(t) - \vec{r}_i(t - \Delta t) + \frac{\Delta t^2}{m_i} \vec{F}_i(t)$$

(repeat ***one billion times*** = microsecond)

Our Microscope is a “Tool to Think”

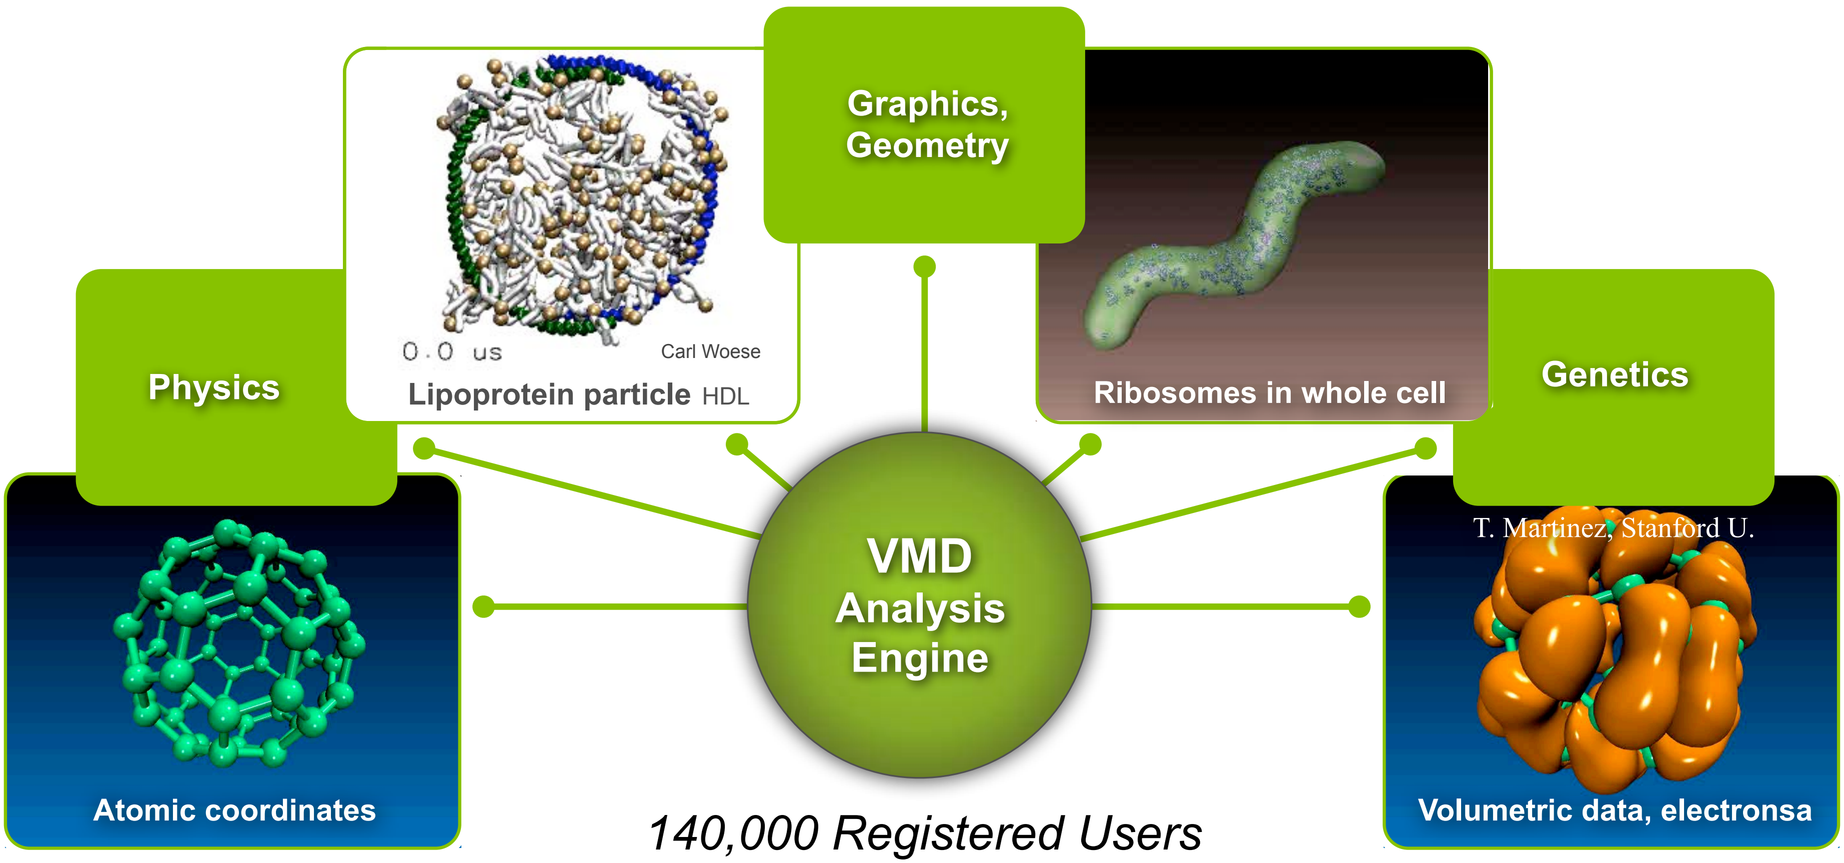

Science 1: Poliovirus Infection

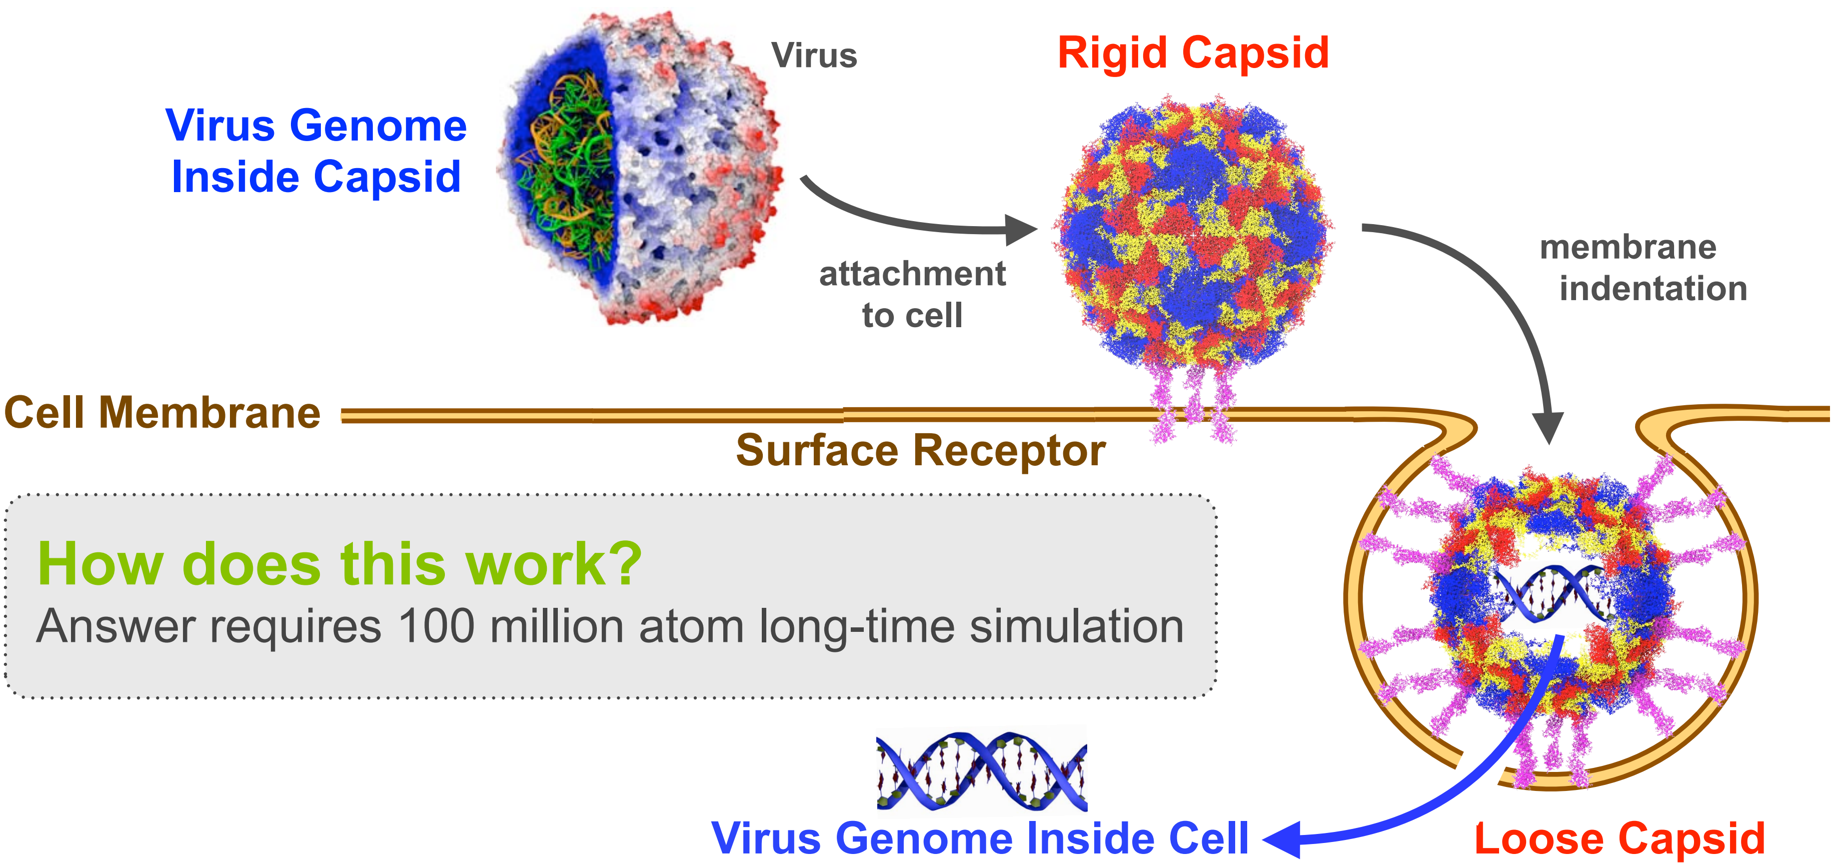

How does this work?

Answer requires 100 million atom long-time simulation

Science 1: Virus Capsid Mechanics

Atomic Force Microscope

— Hepatitis B Virus —

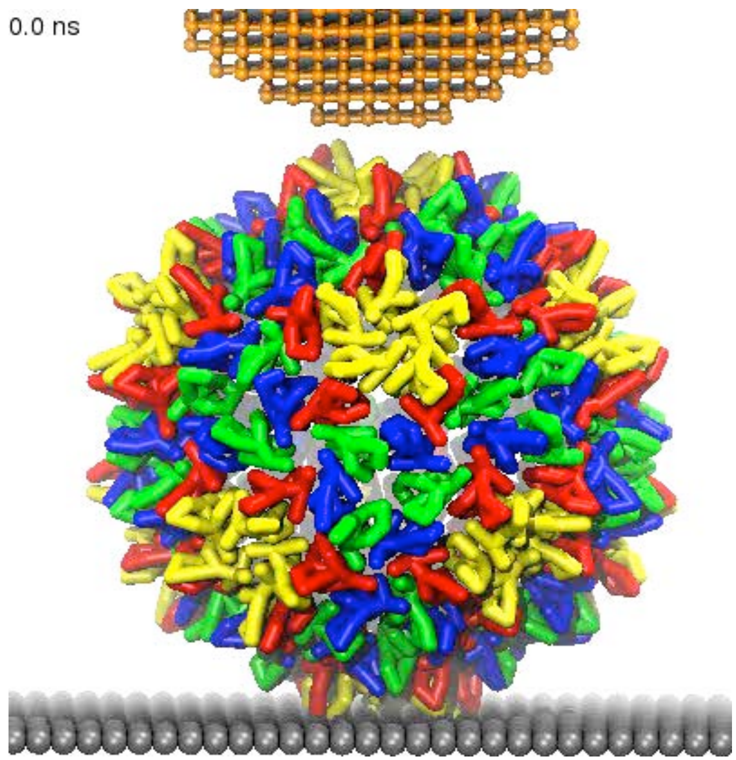

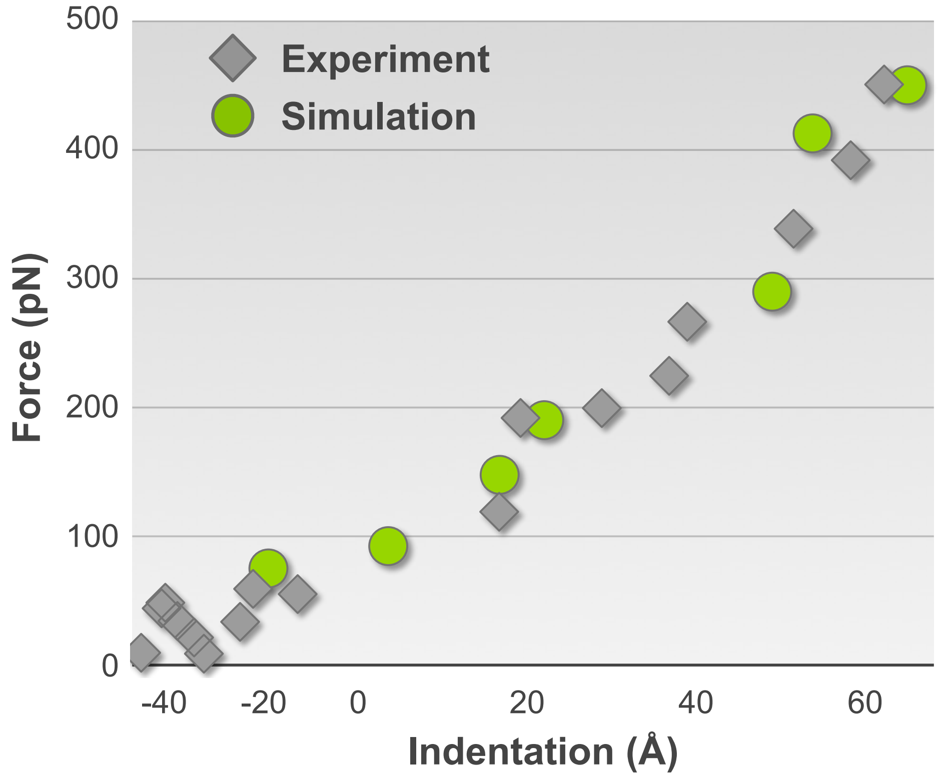

GPU Solution 1: Time-Averaged Electrostatics

- Thousands of trajectory frames
- **1.5 hour** job reduced to **3 min**
- GPU Speedup: **25.5x**
- Per-node power consumption on NCSA GPU cluster:
  - CPUs-only: 448 Watt-hours
  - CPUs+GPUs: 43 Watt-hours
- Power efficiency gain: **10x**

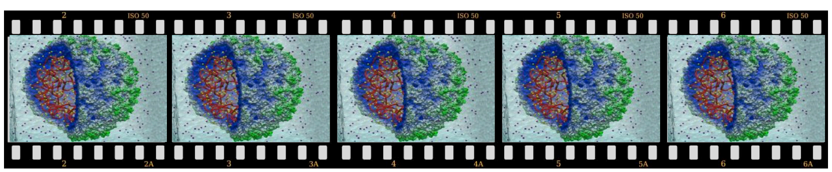

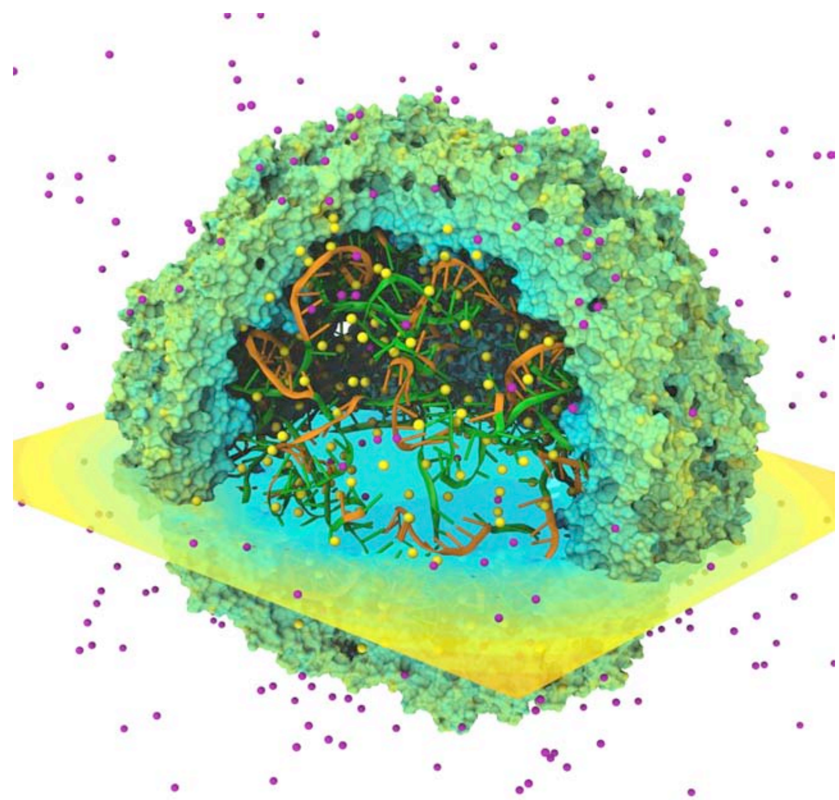

Science 2: How Nature Harvests Sun Light

95% of the energy in the biosphere comes from this energy source

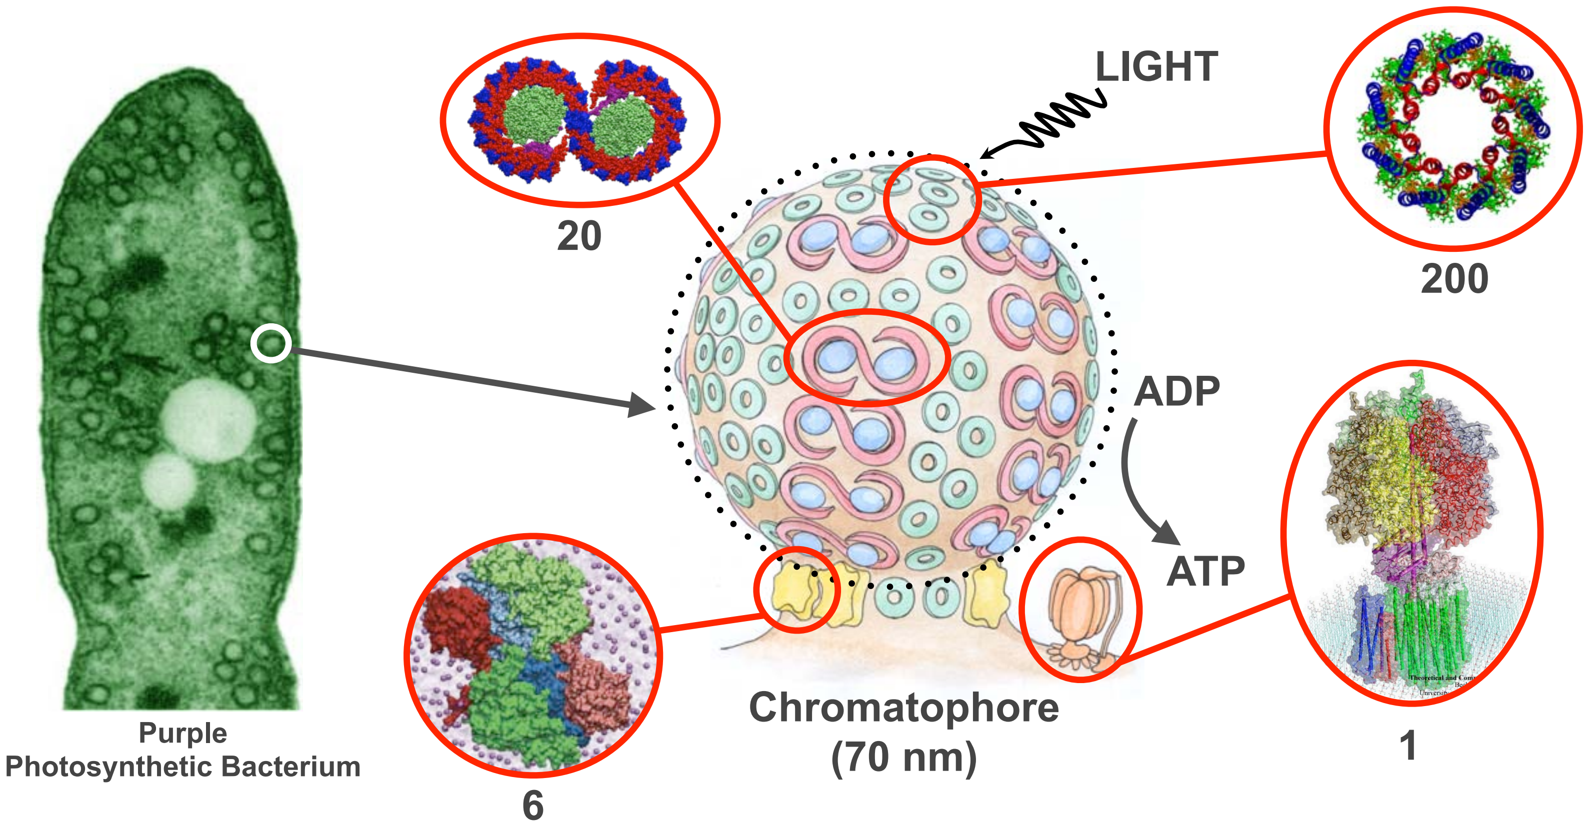

Science 2: How Nature Harvests Sun Light

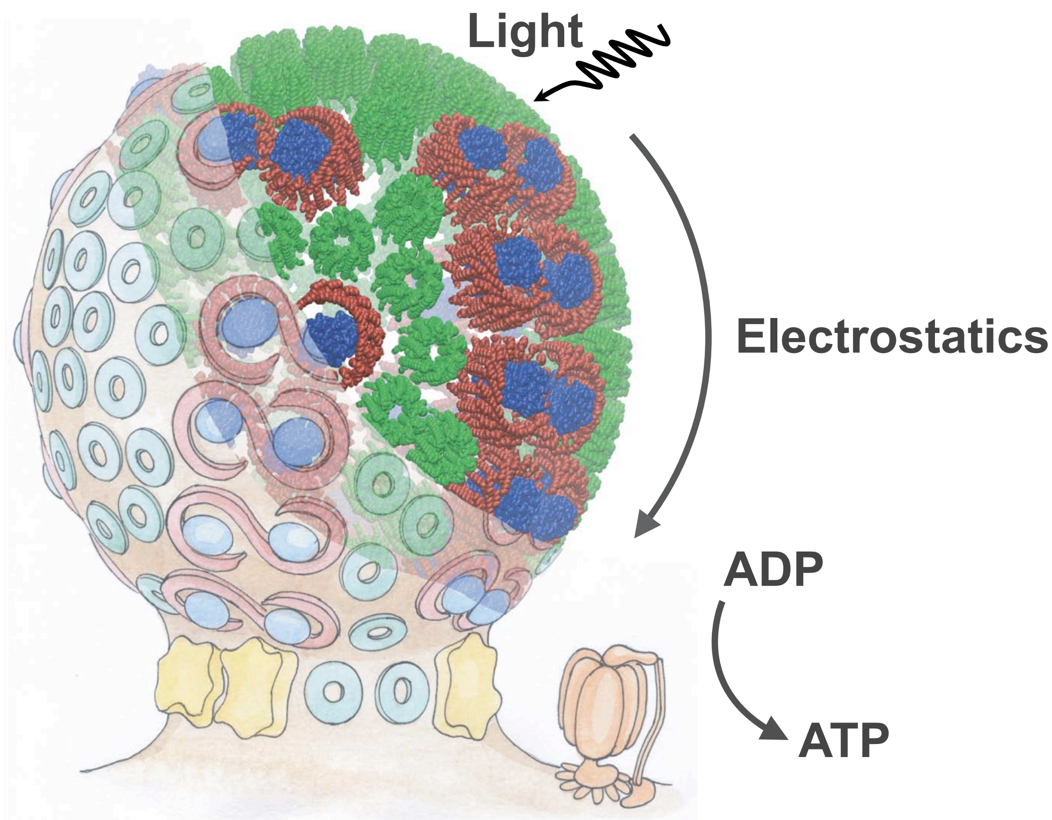

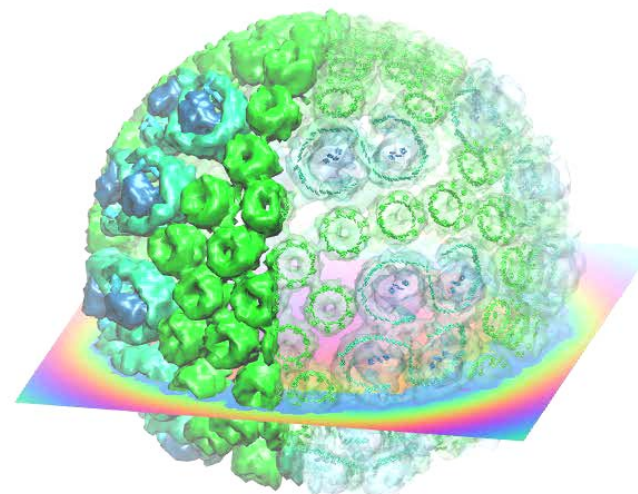

~10M atoms

Electrostatic field calculated
with multilevel summation method

1 CPU core: 1 hr 10 min
3 GPUs (G80): ~90 seconds

GPU Solution 2: Multilevel Summation Method for Electrostatics on the GPU

Localized grid operations map well to GPUs

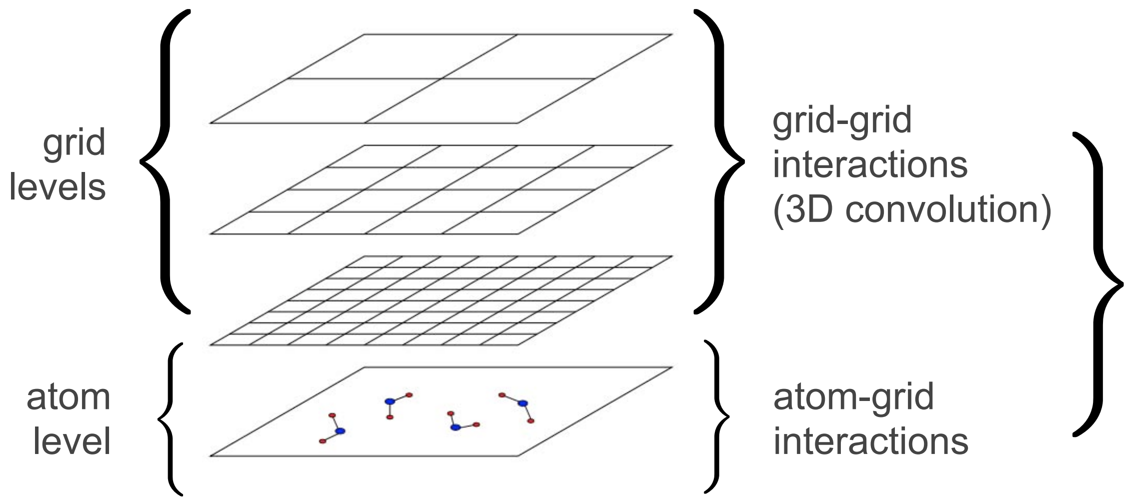

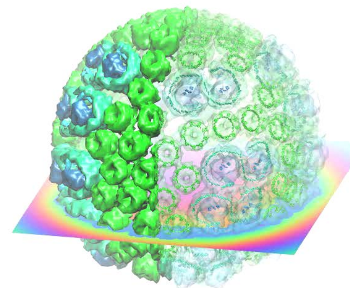

GPU Speedups

26x on 1 GPU
46x on 3 GPUs

10 Million Atoms

1 CPU core: 1 hr 10 min
3 GPUs (G80): ~90 seconds

Multilevel summation method has linear time complexity well suited for GPUs; more flexible than other methods

Science 3: How Proteins are Made from Genetic Blueprint

- Ribosome — Decodes genetic information from mRNA
- Important target of many antibiotics
- Static structures of crystal forms led to 2009 Nobel Prize
- **But one needs structures of ribosomes in action!**

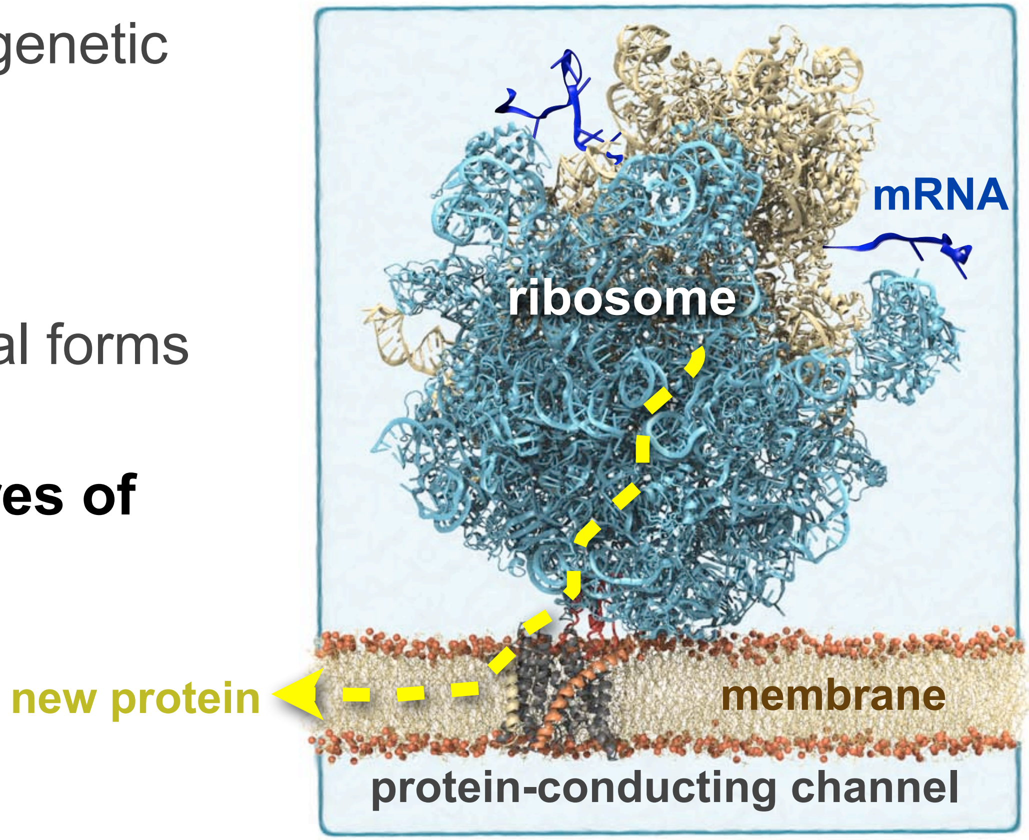

Science 3: How Proteins Are Made from Genetic Blueprint

Low-resolution Data

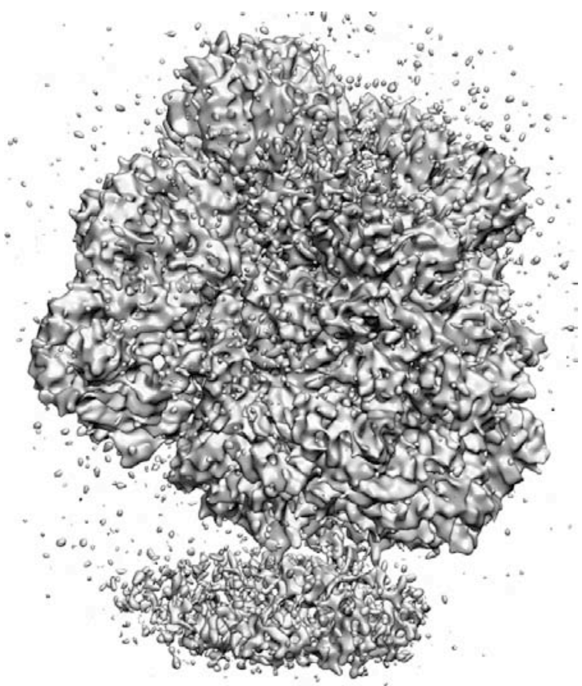

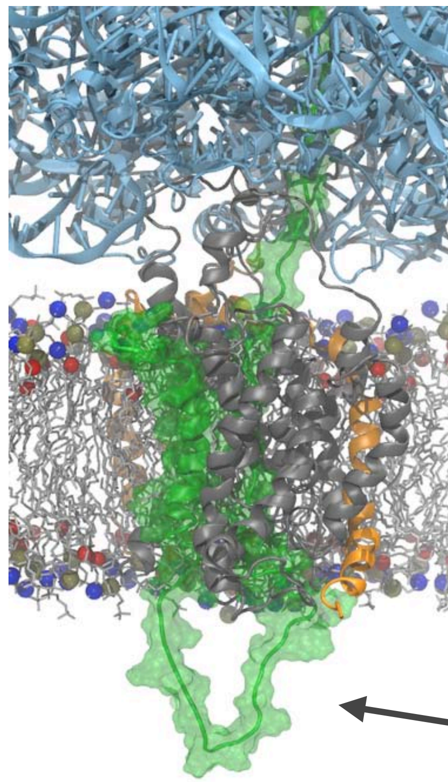

High-resolution Structure

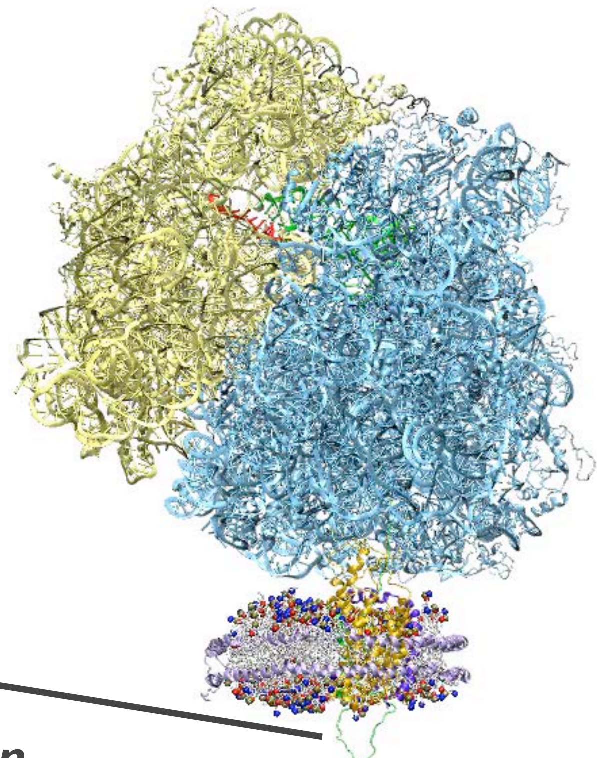

Close-up of Nascent Protein

GPU Solution 3: Molecular Dynamics Simulations

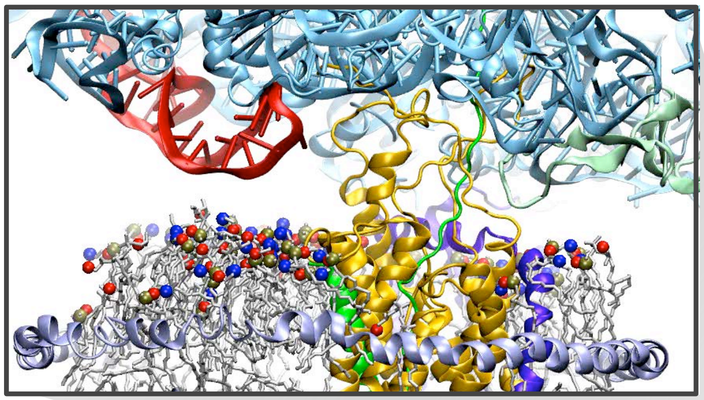

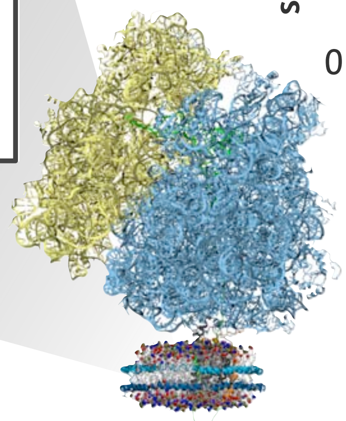

Molecular dynamics simulation of protein insertion process

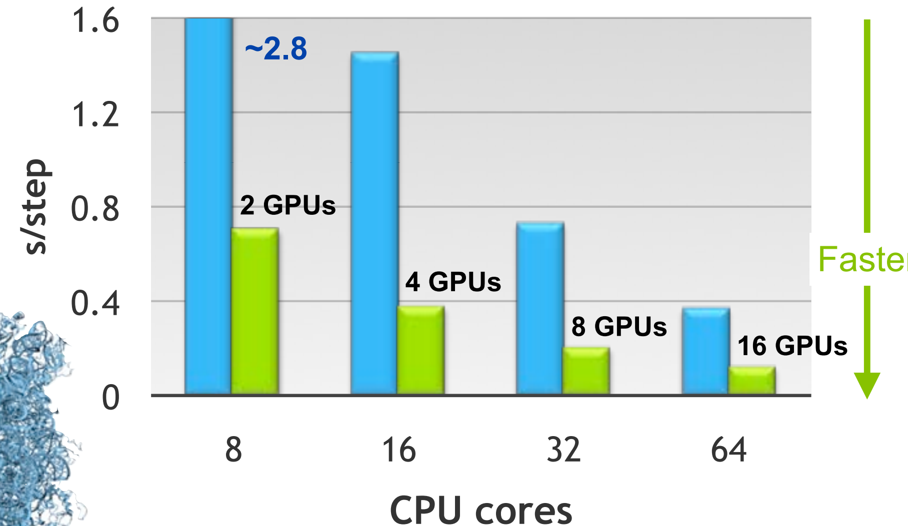

NCSA Lincoln Cluster performance (8 Intel cores and 2 NVIDIA Tesla GPUs per node, 1 million atoms)

GPUs reduced time for simulation from **two months to two weeks!**

Science 4: Nanopore Sensors

Genetics: Genes control our bodies and experiences!
Epigenetics: Our bodies and experiences control the genes!

Epigenetics made possible through DNA methylation

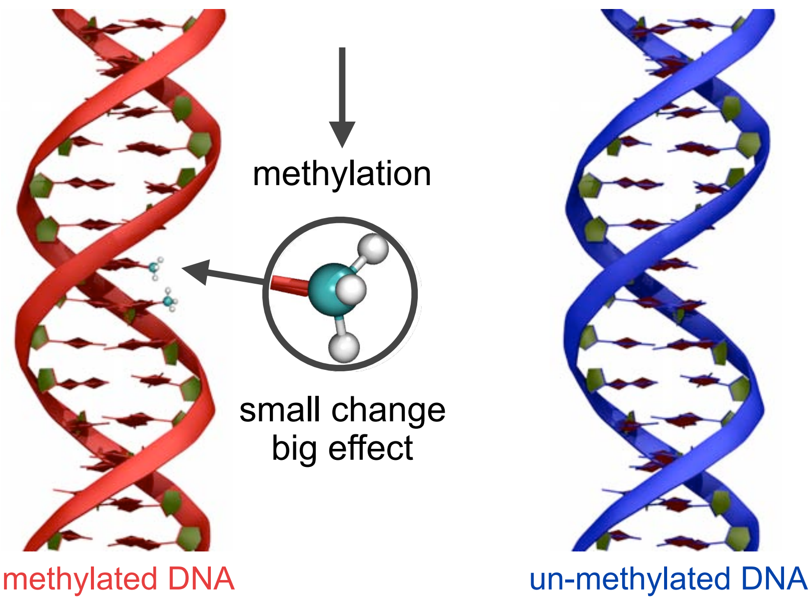

Related pathologies: obesity, depression, cancer

Detect methylation with **nanopores**

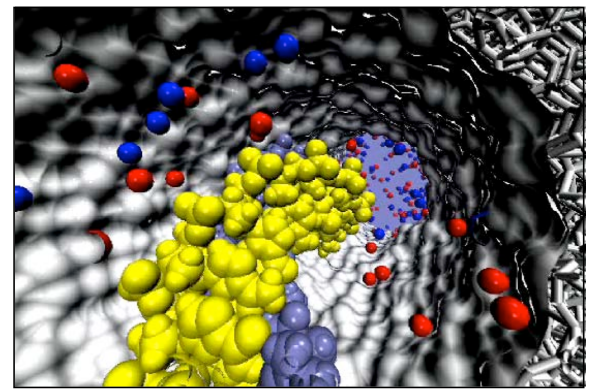

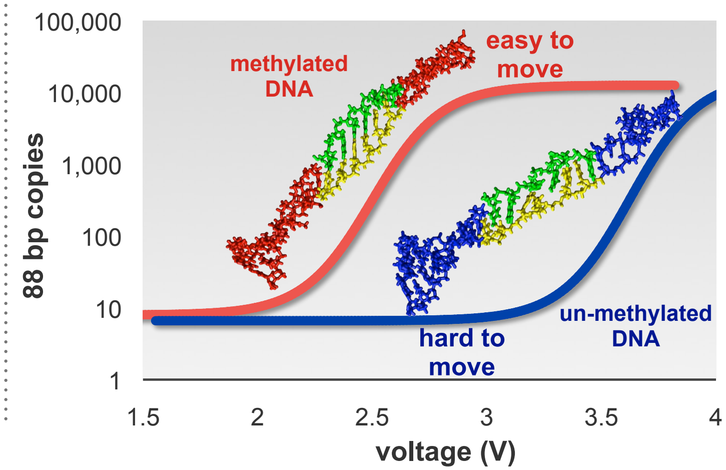

Science 4: Nanopore Sensors

Create a **Better Nanopore** with Polymeric Materials

New materials, new problems: Nanoprecipitation

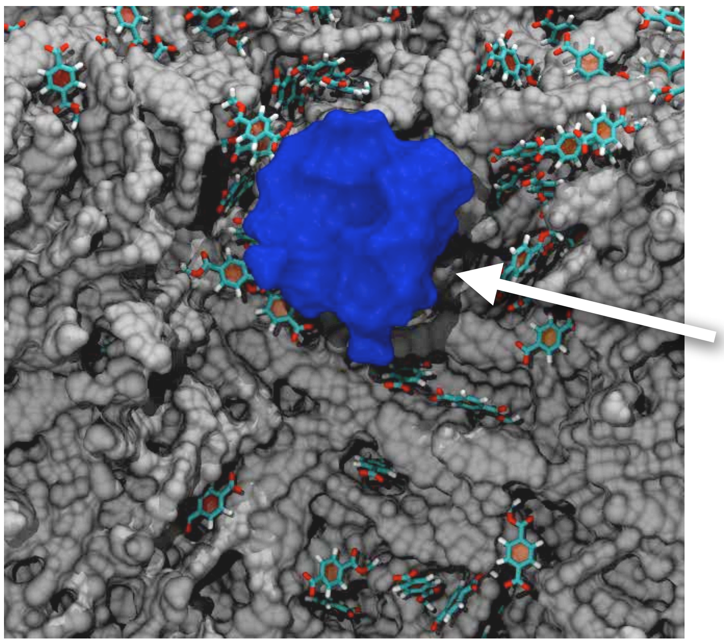

nanoprecipitation of ions

Radial distribution functions
identify nanoprecipitation

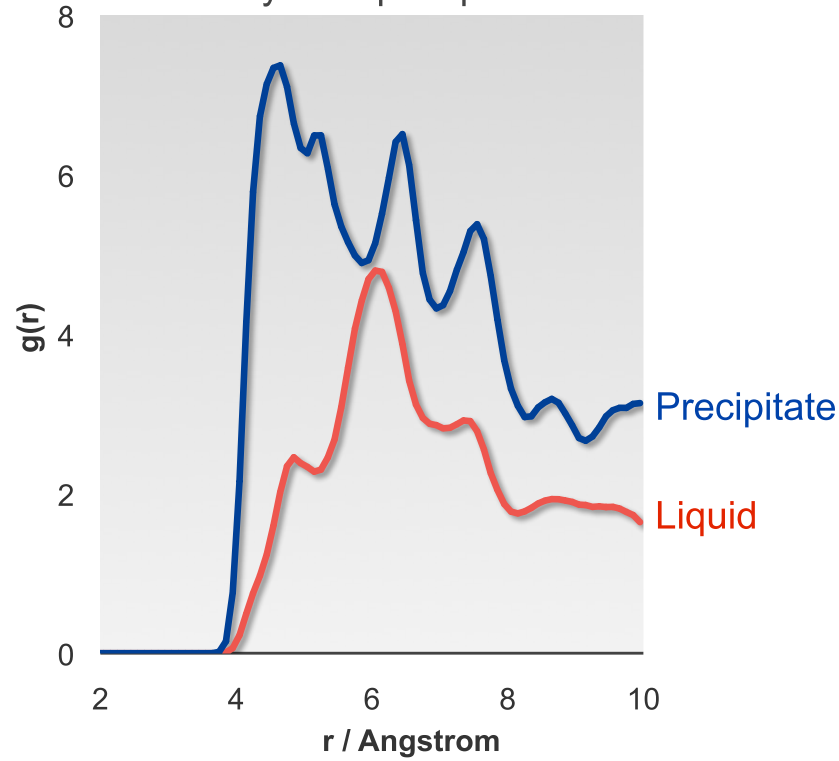

GPU Solution 4: Computing Radial Distribution Functions

- 4.7 million atoms
- 4-core Intel X5550 CPU: **15 hours**
- 4 NVIDIA C2050 GPUs: **10 minutes**
- Fermi GPUs ~3x faster than GT200 GPUs: larger on-chip shared memory

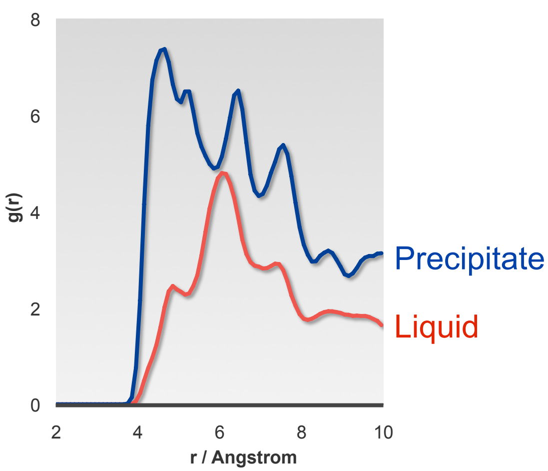

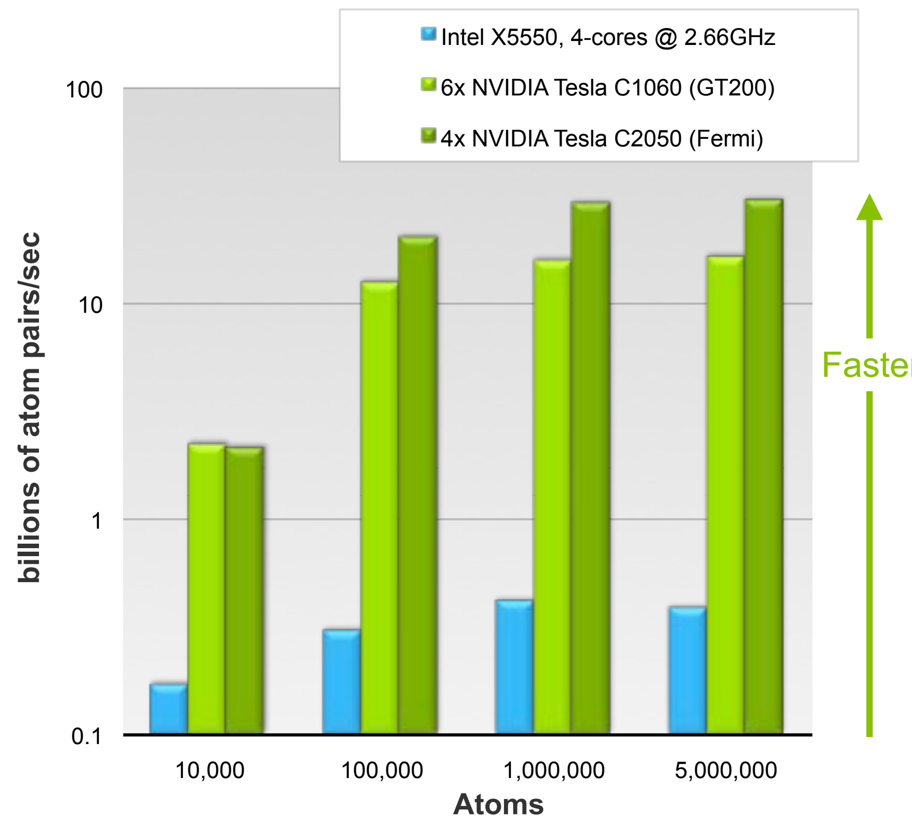

Science 5: Quantum Chemistry Visualization

- Chemistry is the result of atoms sharing electrons
- Electrons occupy “clouds” in the space around atoms
- Calculations for visualizing these “clouds” are costly: **tens to hundreds of seconds** on CPUs – **non-interactive**
- GPUs enable the dynamics of electronic structures to be animated **interactively** for the first time

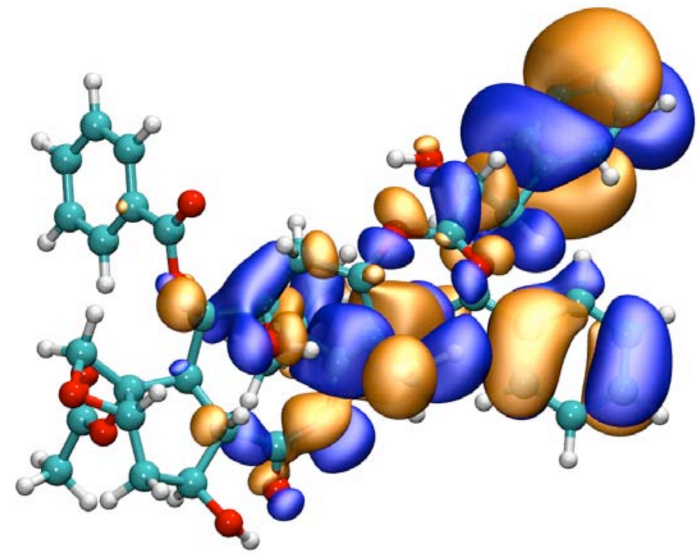

Taxol: Cancer Drug

VMD enables interactive display of QM simulations, e.g. Terachem, GAMESS

Science 5: Quantum Chemistry Visualization

Rendering of electron “clouds” achieved on GPUs as quickly as you see this movie! CPUs: One working day!

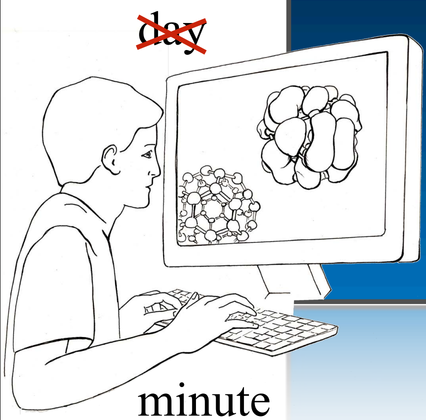

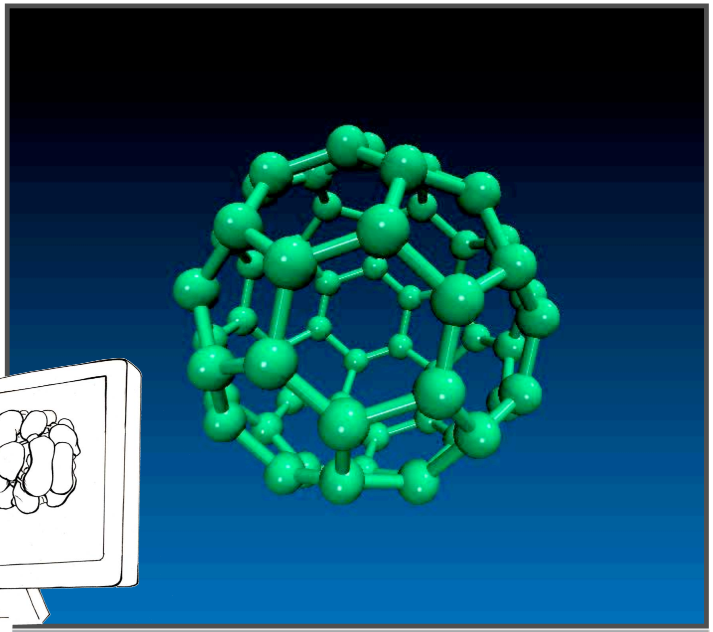

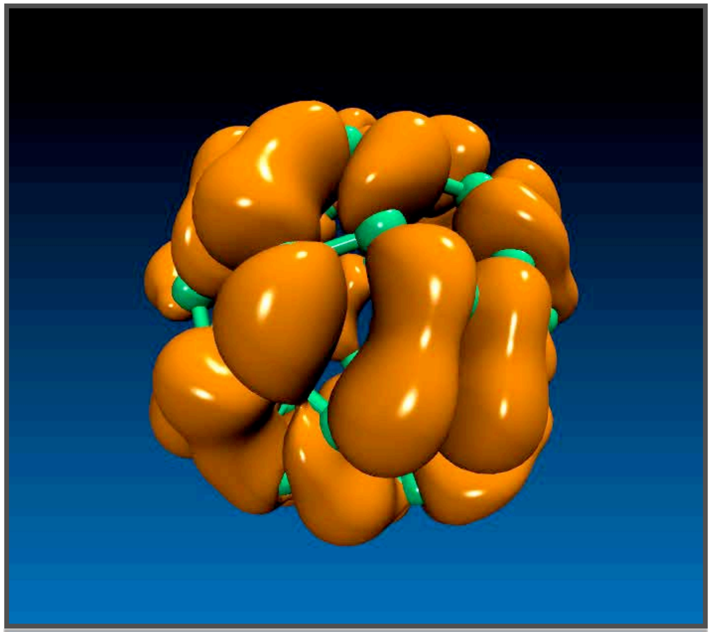

Simulation: Terachem Interactive Visualization: VMD
Courtesy T. Martinez, Stanford

GPU Solution 5: Computing C_{60} Molecular Orbitals

Device	CPUs, GPUs	Runtime (s)	Speedup
Intel X5550-SSE	1	30.64	1.0
Intel X5550-SSE	8	4.13	7.4
GeForce GTX 480	1	0.255	120
GeForce GTX 480	4	0.081	378

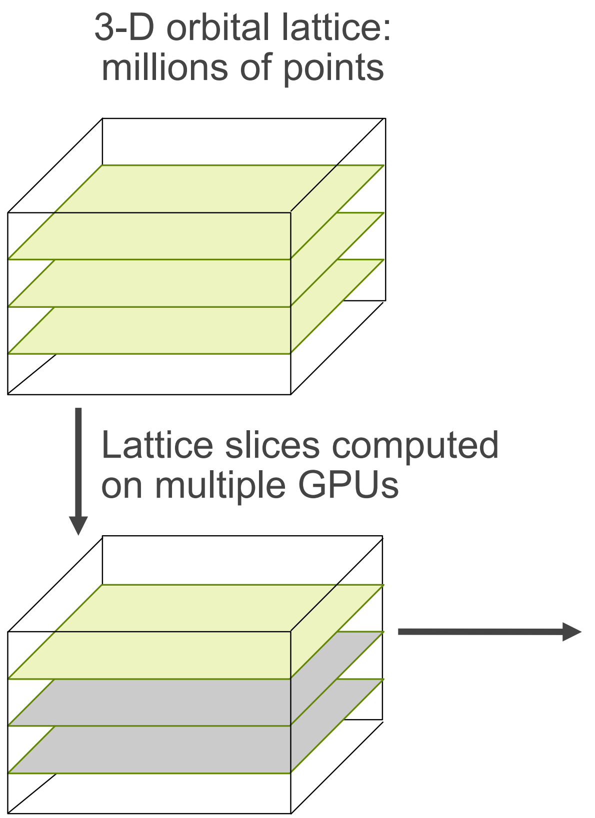

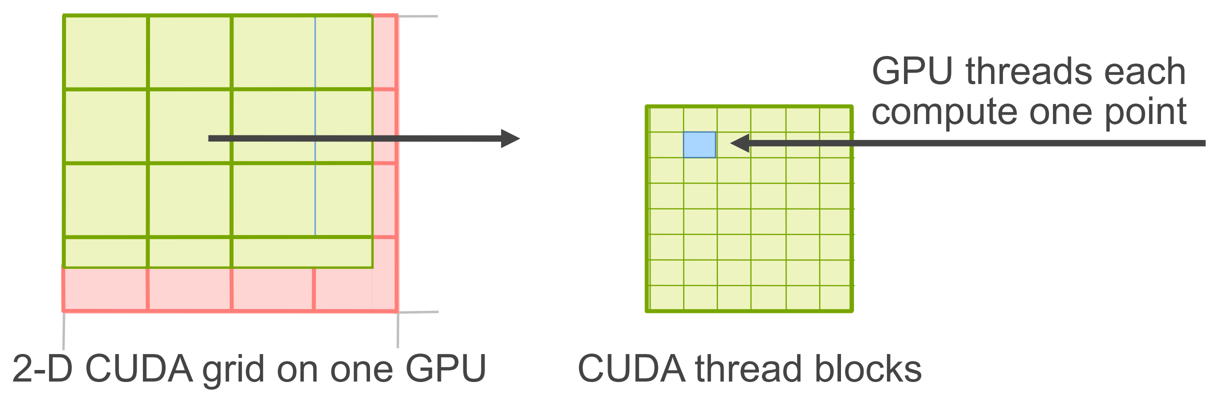

Science 6: Protein Folding

- Protein **misfolding** responsible for diseases:
  - Alzheimer's
  - Parkinson's
  - Huntington
  - Mad cow
  - Type II diabetes
  - ...

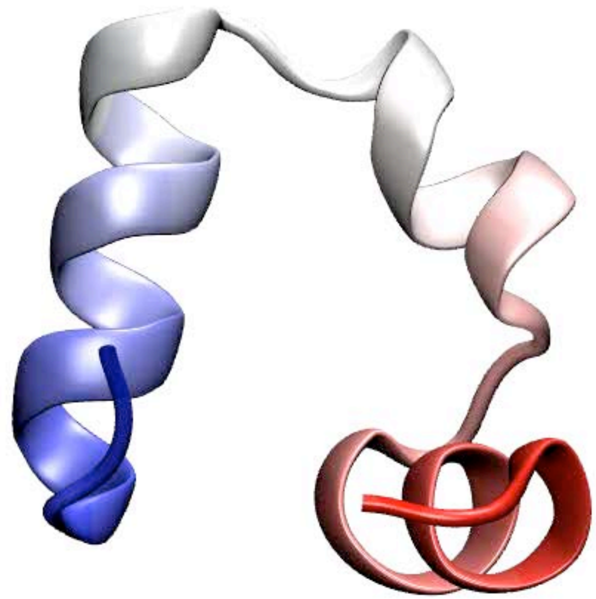

villin headpiece
3 months on 329 CPUs

Observe folding process in unprecedented detail

Science 6: Protein Folding

- Some simulations still fail to fold proteins due to inaccurate modeling of interatomic forces!
- Protein folding demands more **accurate** model which leads to more **expensive** computation

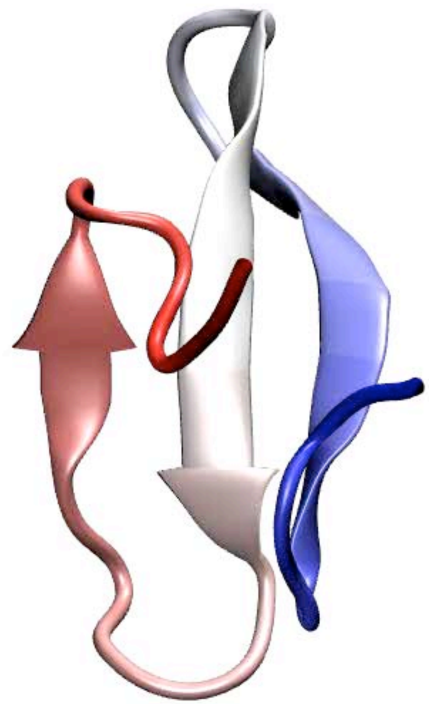

WW domain
3 months on 329 CPUs

GPU Solution 6: Computing More Accurate Simulation Models

- Atomic polarizability increases computation by 2x…
- …but, the additional computations are perfectly suited to the GPU!
- For now, NAMD calculates atomic polarizability on CPUs only...soon we will also use GPUs

NAMD CPU performance scaling

polarizable water
non-polarizable water

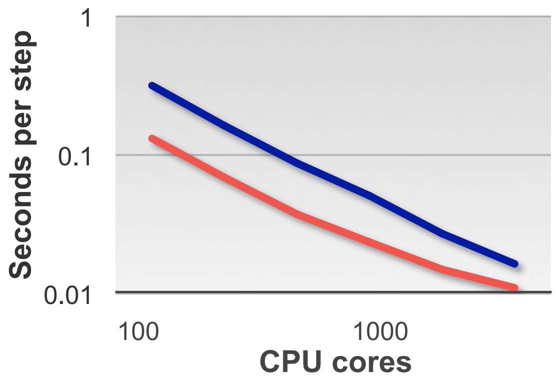

Atomic polarizability of water, highly accurately simulated through additional particles (shown in green)

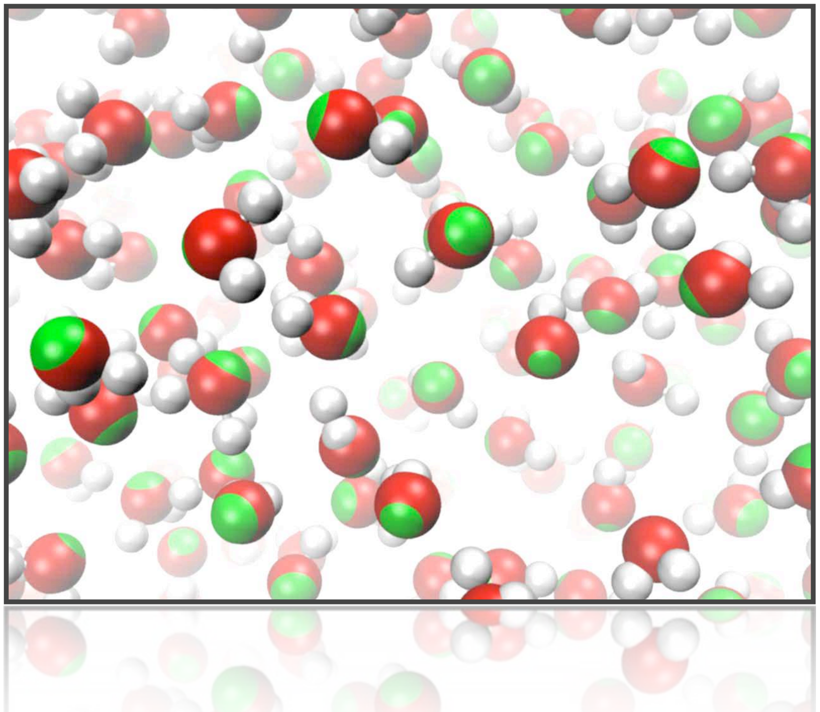

Genetic activity of E. coli bacteria

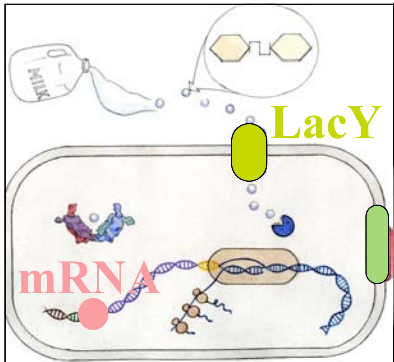

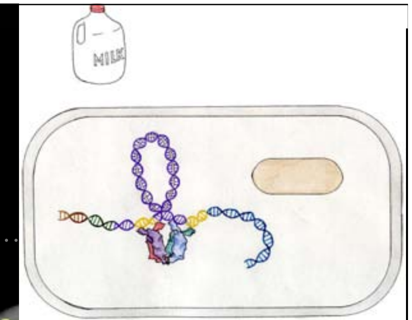

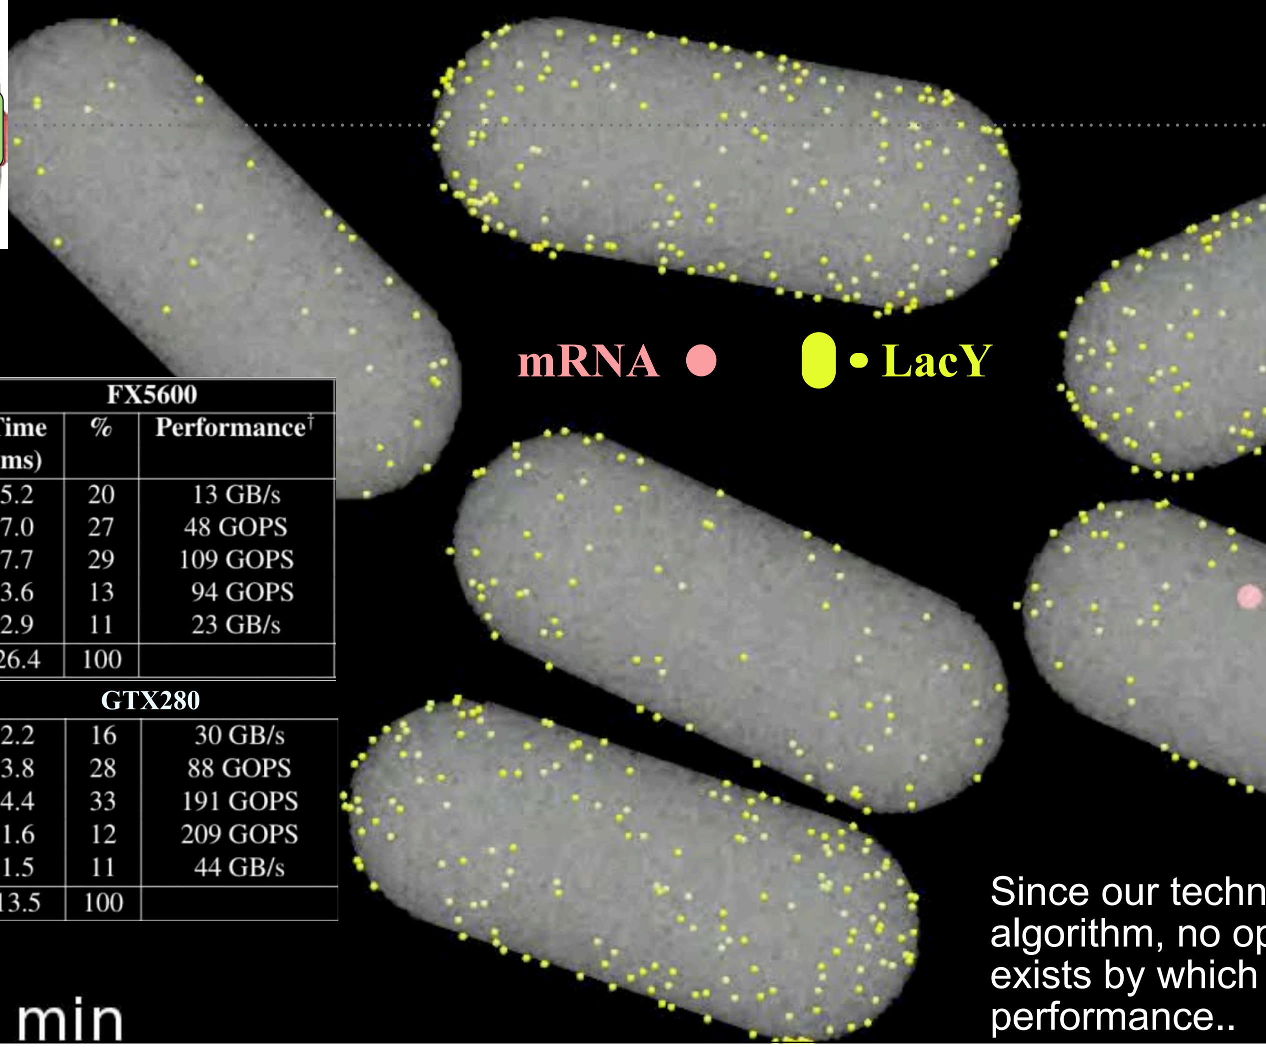

Calculation	FX5600		
	Time (ms)	%	Performance†
Load lattice block	5.2	20	13 GB/s
Random number generation‡	7.0	27	48 GOPS
Particle movement decision	7.7	29	109 GOPS
Particle propagation	3.6	13	94 GOPS
Store lattice block	2.9	11	23 GB/s
Total	26.4	100	
	GTX280		
Load lattice block	2.2	16	30 GB/s
Random number generation‡	3.8	28	88 GOPS
Particle movement decision	4.4	33	191 GOPS
Particle propagation	1.6	12	209 GOPS
Store lattice block	1.5	11	44 GB/s
Total	13.5	100	

Zan Luthey-Schulten
and Elijah Roberts

Since our technique is a native GPU algorithm, no optimized CPU version exists by which to measure its performance..

2010 Workshop on GPU Computing for Molecular Modeling

- Spread the benefits of GPU computing to solve new problems in molecular modeling
- Intensive 2-day workshop after 1-week GPU workshop at NCSA
- Participants present their work and exchange ideas and GPU solutions

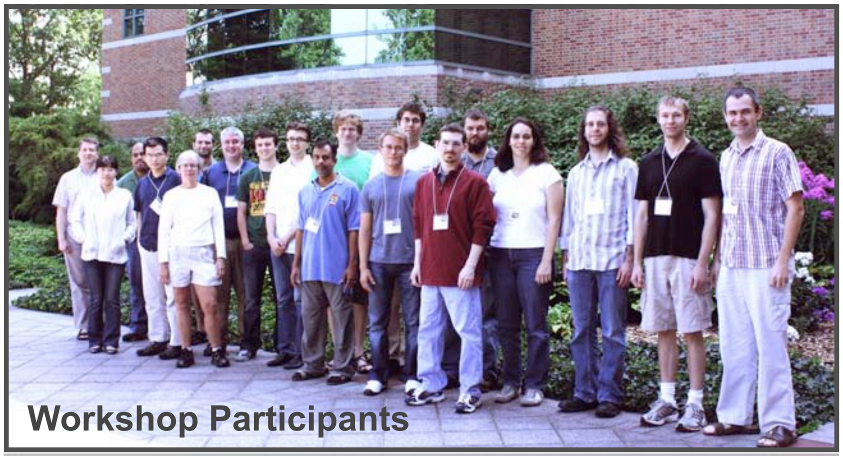

Workshop Participants

Three of our GPU Heroes

Our GPU Biomedical Science Computing Goals:

- More accurate simulations
- Speed-up: simulations now take minutes instead of weeks
- Make previously unreachable scales accessible

Acknowledgements

- Theoretical and Computational Biophysics Group, University of Illinois at Urbana-Champaign
- Wen-mei Hwu and the IMPACT group at University of Illinois at Urbana-Champaign
- L. Kale and the Center for Parallel Computing at University of Illinois at Urbana-Champaign
- NVIDIA CUDA Center of Excellence, University of Illinois at Urbana-Champaign
- Ben Levine, Axel Kohlmeyer at Temple University
- NCSA Innovative Systems Lab
- The CUDA team at NVIDIA
- Zan Luthey-Schulten and Elija Roberts (E. coli whole cell simulations)
- NIH support: P41-RR05969

GPU TECHNOLOGY CONFERENCE

Thank You

PRESENTED BY NVIDIA.