Efficient Multiscale Platelets Modeling using Supercomputers

Na Zhang
Advisor: Professor Yuefan Deng
Department of Applied Mathematics and Statistics
Stony Brook University

Motivation

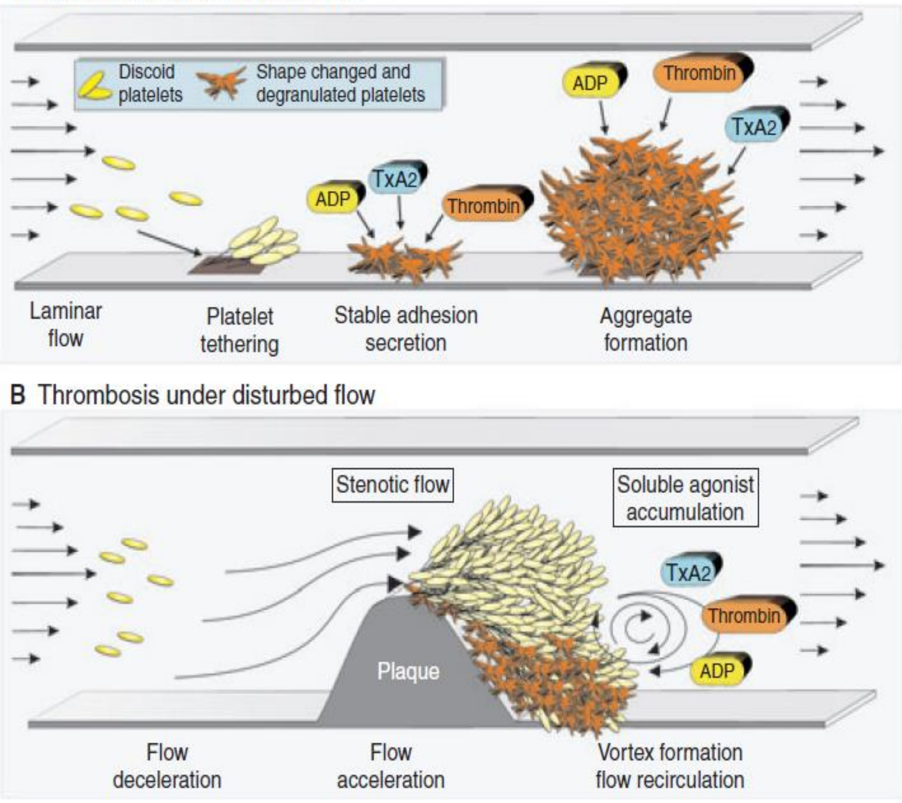

Jackson et al., “Dynamics of Platelet Thrombus Formation,” Journal of Thrombus and Haemostasis”, 2009

Activation of Platelets in Cardiovascular Devices

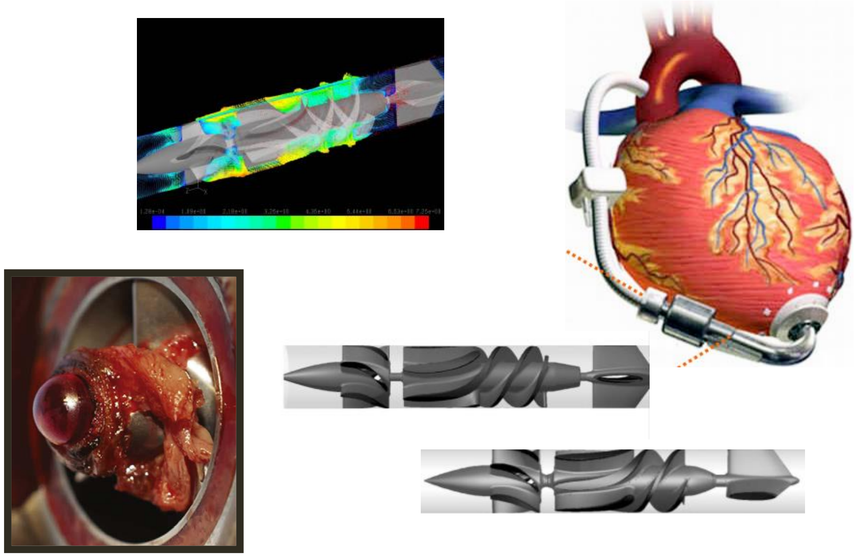

CFD and VADs pictures are from: Girdhar, G., Xenos, M., *et al.*, “Device thrombogenicity emulation: a novel method for optimizing mechanical circulatory support device thromboresistance”, 2012. Thrombosis picture is from: Mokadam, N.A. *et al.*, “Thrombus formation in a HeartMate II”, 2011.

Project: Multiscale Modeling Of Blood Flow and Clotting In Cardiovascular Devices

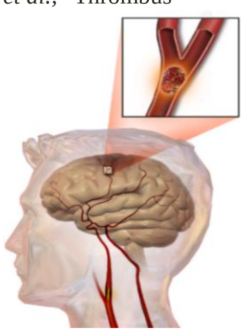

Design of Large Scale Simulations – Integration of Multiscale Model, Numerical Algorithm, and Supercomputing

Simulation Method Considerations

- **Multiscale Modeling with Particle-Based Methods**
  - **Mesoscale: Dissipative Particle Dynamics (DPD)-Based fluid-modeling**
  - **Microscale: Coarse-Grained Molecular Dynamics (CGMD)-Based platelet-modeling**
- **Why not using Computational Fluid Dynamics (CFD) or Lattice Boltzmann Method (LBM)?**
  - Fail to capture small-scale molecular mechanisms upon platelet shape change and interactions of key players in blood coagulation
  - Hard to do interfacing/coupling
- **Why not using Quantum or Fine-Grained Molecular Dynamics(MD)?**
  - Computationally infeasible

Scale	Microscale	Mesoscale
Approach	CGMD	DPD
Domain	Platelet	Blood Plasma
Time	**10 ~ 100 fs**	**0.01 ~ 1 μs**
Space	**1 ~ 20 Å**	**0.1 ~ 1 μm**
Model	Particle-Based	

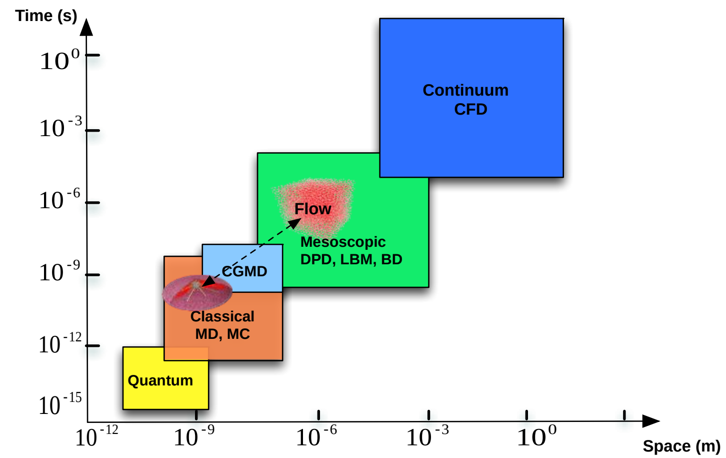

Schematic representation of platelet and viscous flow at the multiple spatial and temporal scales

Multiscale Model Overview

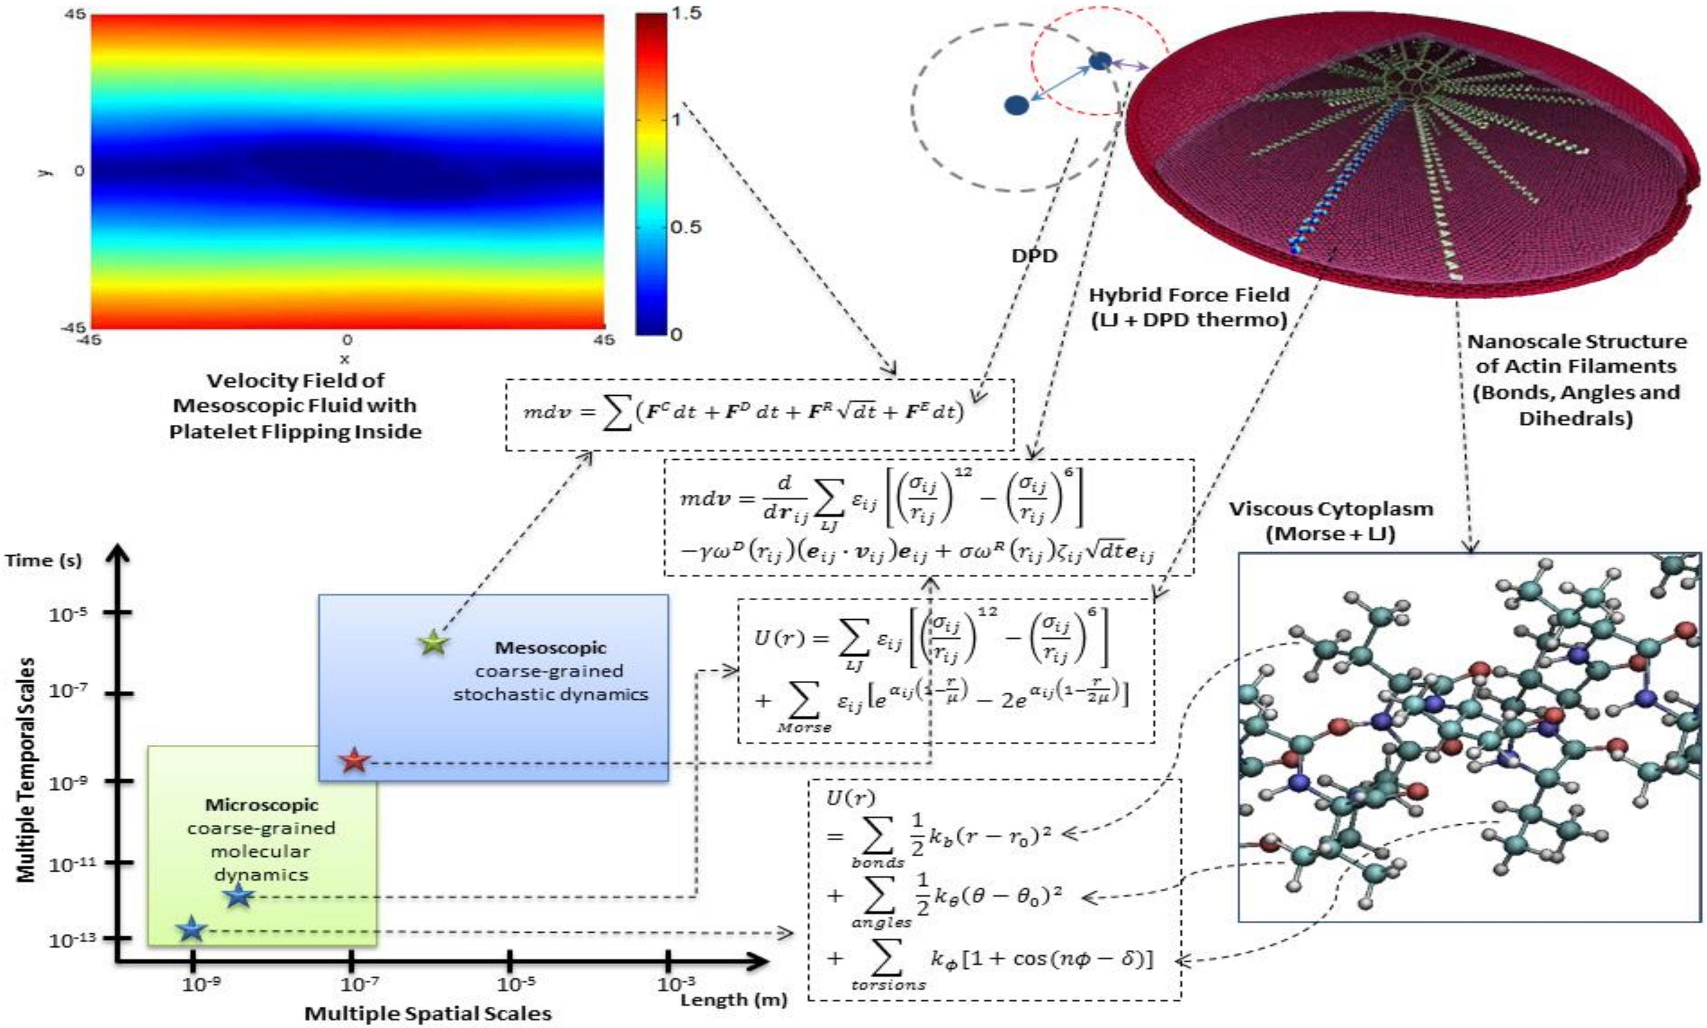

Interfacing Nano-Mesoscopic Blood Flow Induced Platelet Activation

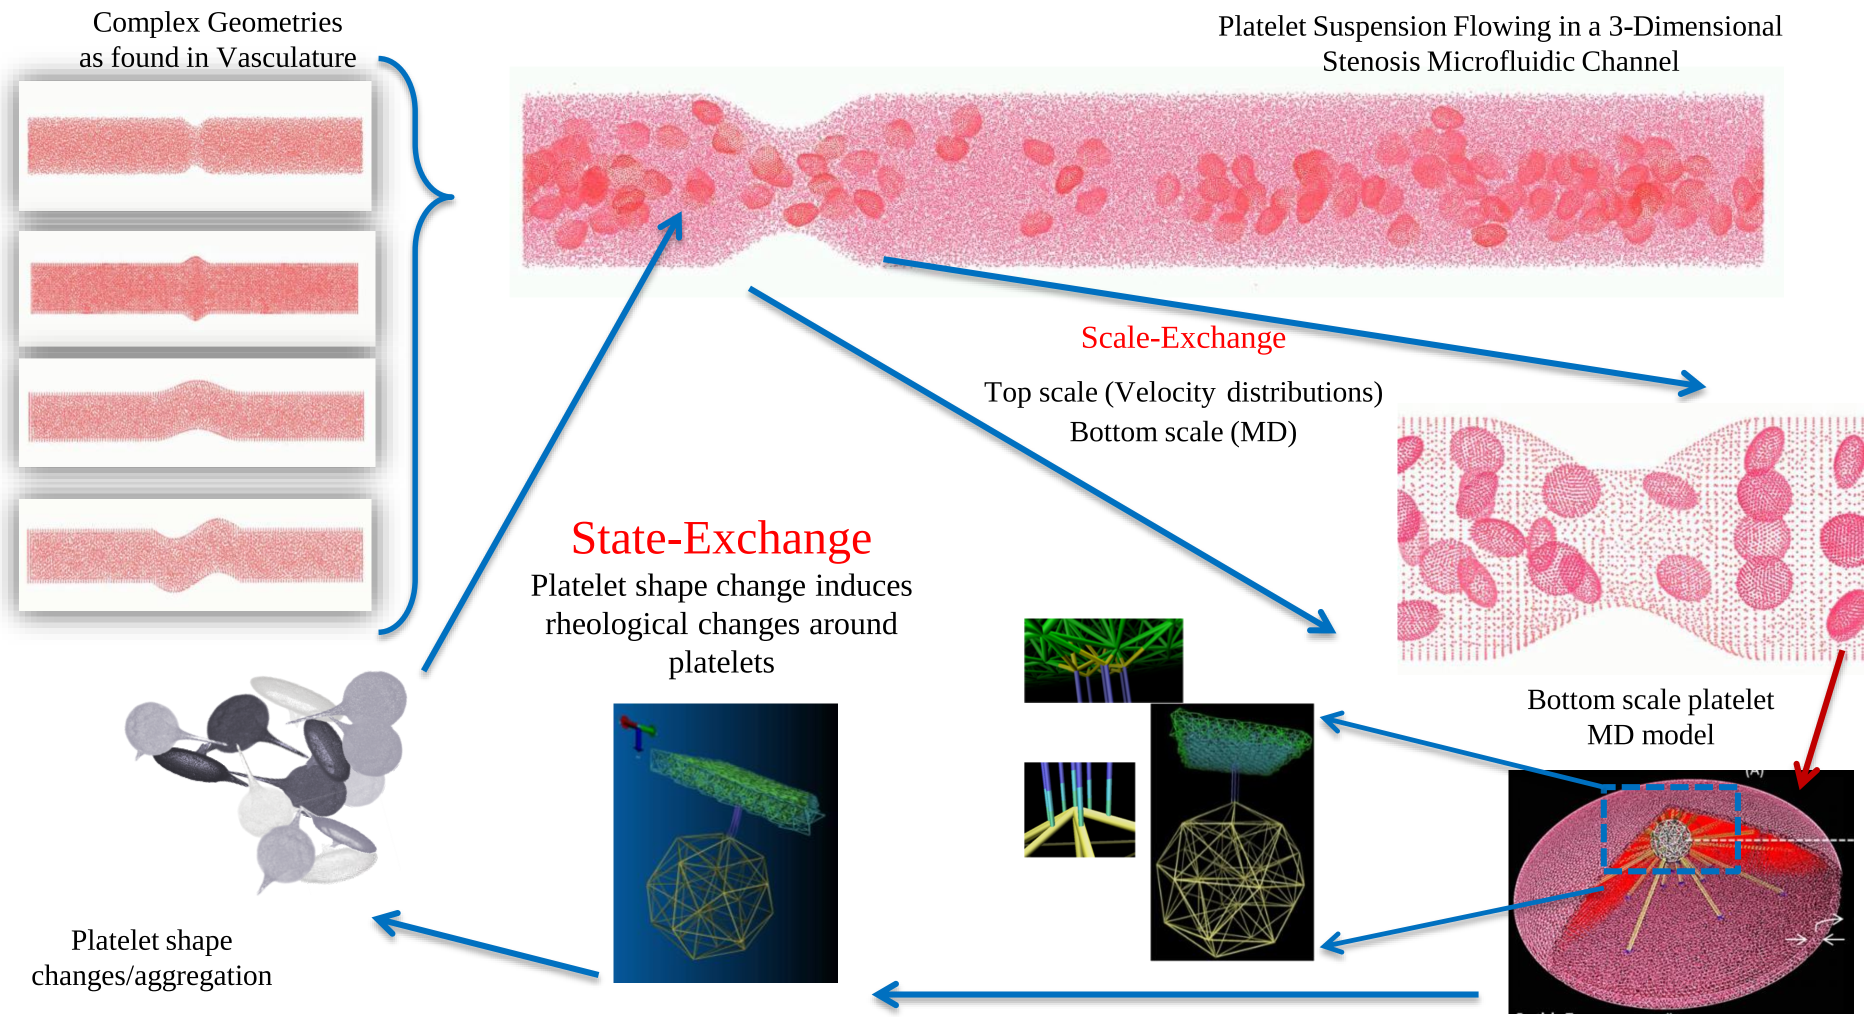

Platelet shape change:
Filopodia protrusion through the
platelet elastic membrane

Fine-Grained Platelet: 147K particles
Two-layer membrane and cytoplasm coupled with
cytoskeletal elements

Speedup Strategies

➢ Computing Challenges and Speedup Strategies:

Categories	Single Platelet	Multiple Platelets
In Vacuum	~0.14 million particles	Complex interactions among platelets
In Blood Plasma	~0.6 million particles	~2.7 million particles for 4 platelets flipping in blood plasma ~10.9 million particles for 16 platelets flipping in blood plasma
In Blood Vessels	Many types of blood cells and complex interactions among those cells and injured walls	
With Shear Stresses & Thermo Conditions	Complex inputs and outputs control; On-the-fly analysis of large datasets	

Bottlenecks	Parallel Computing Strategies
Disparate Time Scales	Multiple Time Stepping Algorithm
Hardware Acceleration	GPU Enabled Computing
Inter-Processor Communication	High-Density Multi-GPU Accelerating

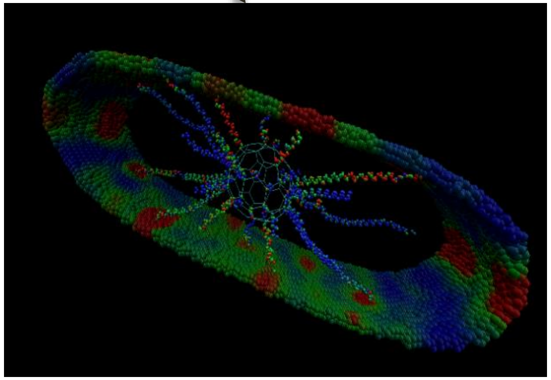

Platelet Surface Stress Field

Multiple Time Stepping Algorithm

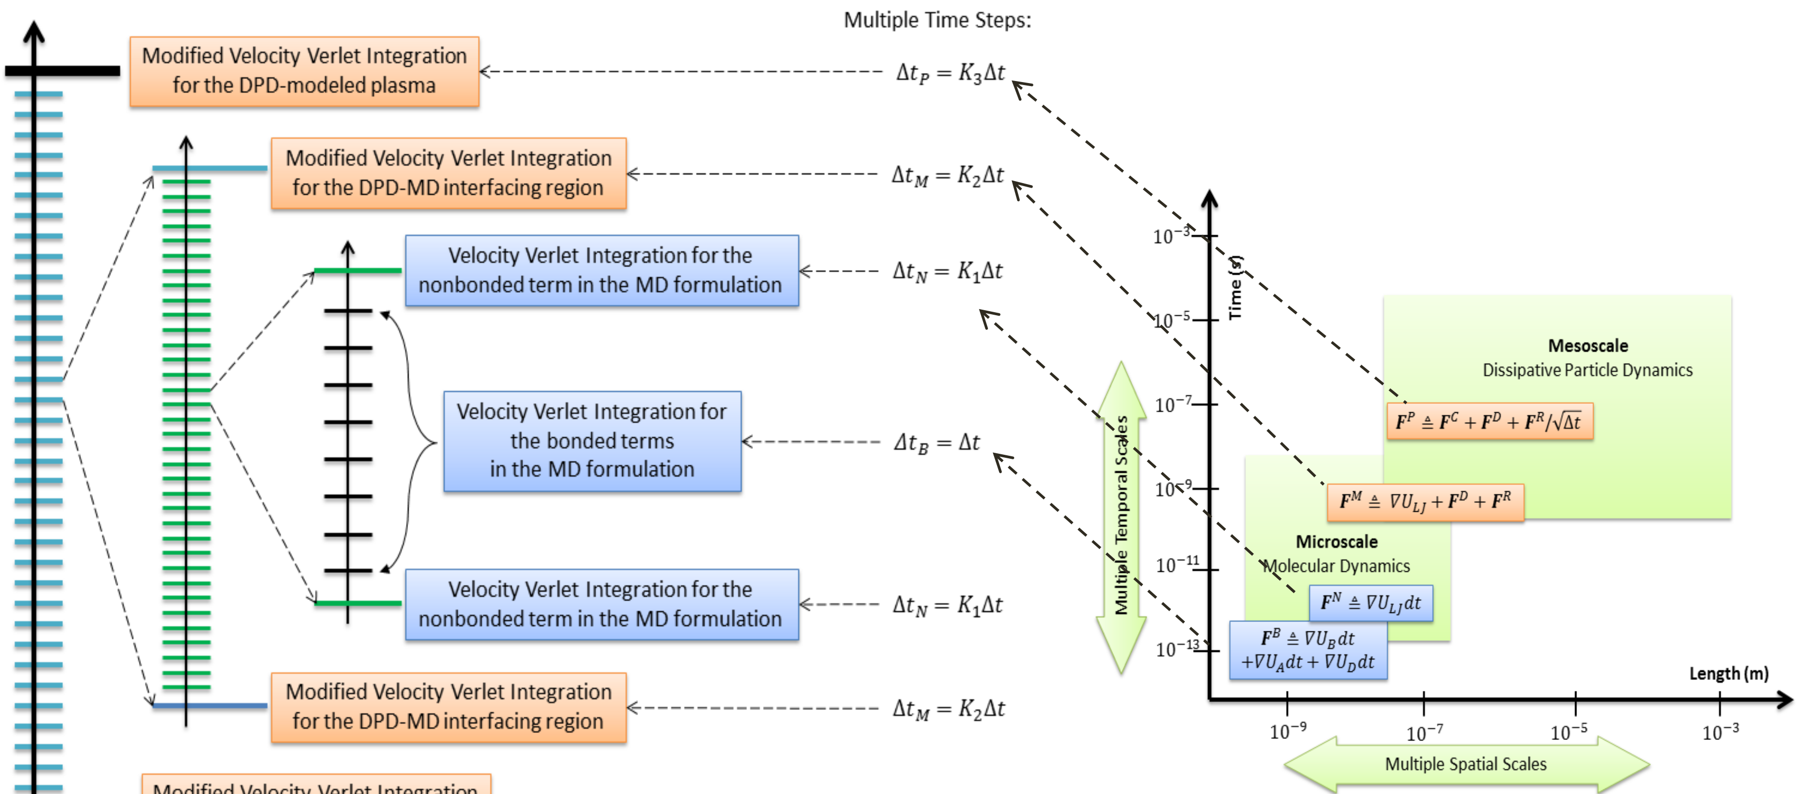

K_1, K_2, K_3 are "Jump Factors"

Multiple Scales in the Model

Results of Multiple Time Stepping Algorithm

➢ Performance Gains by MTS with CPU-only Solutions

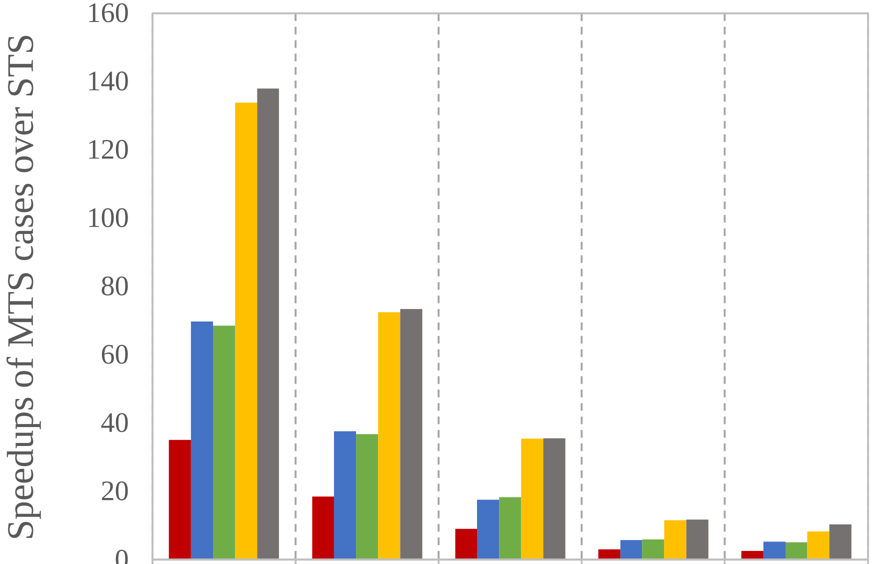

	np32	np64	np128	np256	np512
CaseA	35.04	18.46	8.94	2.96	2.55
CaseB	69.68	37.54	17.47	5.67	5.24
CaseC	68.48	36.75	18.30	5.92	5.05
CaseD	133.86	72.45	35.42	11.54	8.19
CaseE	137.90	73.39	35.48	11.67	10.34

Times faster of MTS algorithms over the STS algorithm

Platform: Stampede
Usage: 32 ~ 512 Procs (2 ~ 32 CPU nodes)
Problem Size: 2.7 million particles (4 platelets flipping in blood plasma)

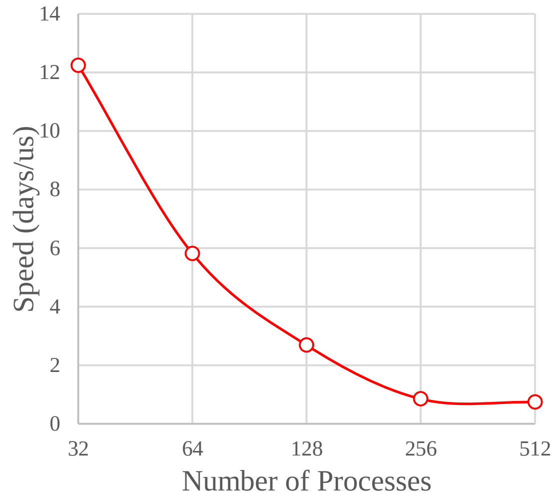

Problem scalability by STS

Stony Brook University

Communication Bottleneck

- Performance Profiling

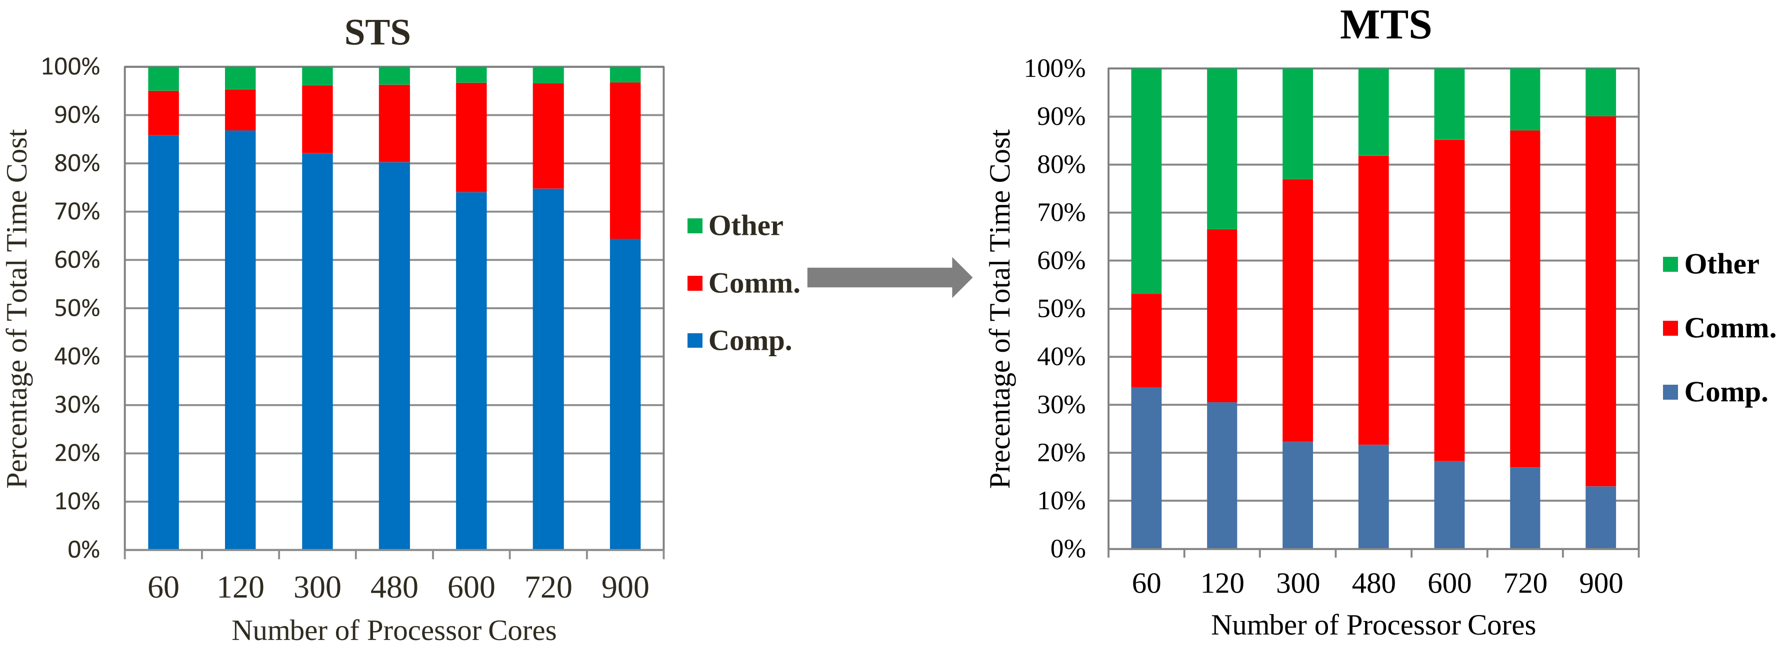

Results of GPGPU Acceleration

➢ Performance Gain by GPGPU Acceleration Based on MTS Gains

- 2~3 times faster over CPU-only for different MTS cases;
- Total speedup of 23 over STS using 16 GPU nodes (K20, Stampede) → Simulating 1-*ms* multi-scale phenomena is reduced to approximate 37 days from 850 days.

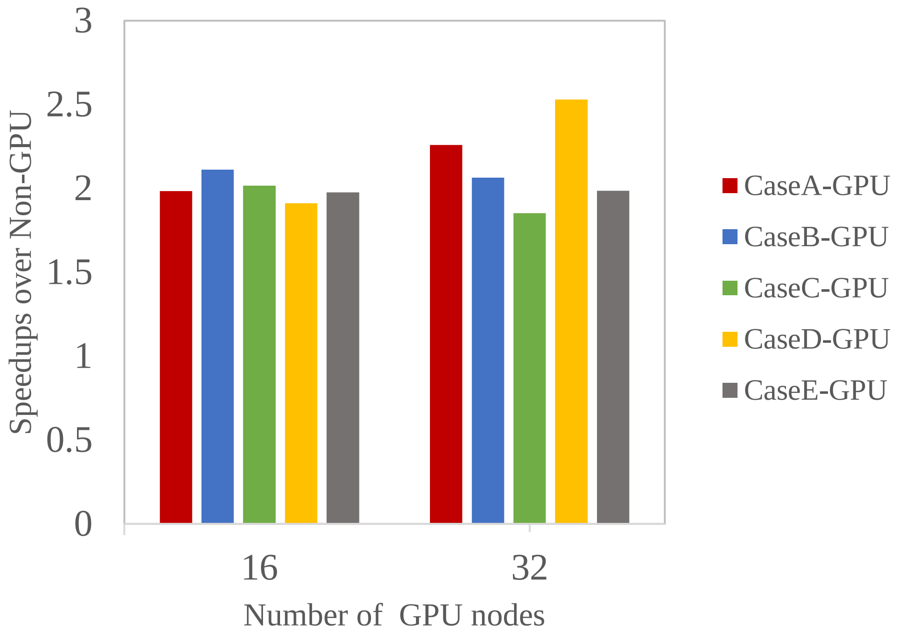

Platform: Stampede
Usage: 16 or 32 GPU nodes
Problem Size: 2.7 million particles
(4 platelets flipping in blood plasma)

Times faster GPU solutions over CPU-only solutions for different MTS cases

Other Performance Examinations

- **Stampede:**
  - Peak Performance: 10 Petaflops
  - 6400 number of compute nodes each with 2 Intel Xeon E5 (Sandy Bridge) processors and an Intel Xeon Phi Coprocessor and some are augmented with a NVIDIA K20 GPU
  - Interconnect: InfiniBand FDR

Stampede

- **Tianhe-2 (Milky Way-2):**
  - Peak Performance: 33.9 Petaflops
  - 3120000 total cores and 2736000 accelerator/co-processor cores
  - Interconnect: TH Express-2

TH-2

- **Cray CS-Storm:**
  - High-density multi-GPU server
  - Up to 8 NVIDIA Tesla K40m GPU devices per node or 16 Tesla K80 GPU devices per node
  - Almost 250 teraflops per rack

Cary CS-Storm Chassis

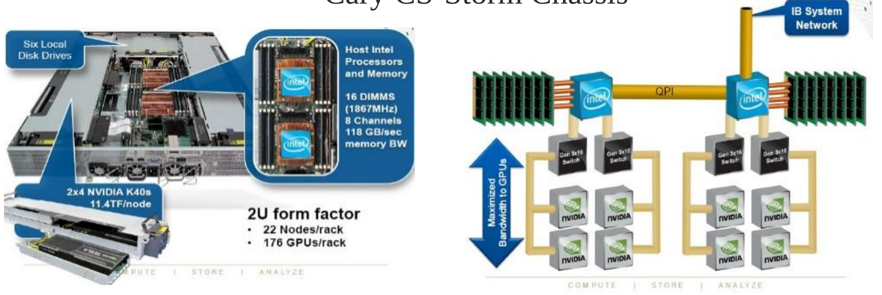

For Results, please refer to my Poster

Source: http://www.cray.com/products/computing/cs-series/cs-storm

Summary and Future Work

- The computational methodology using **multiscale models and algorithms on supercomputers** could offer a promising approach for modeling platelet-related phenomena, in an attempt to better design drugs for fighting vascular diseases.
- The combined acceleration strategy, i.e., **the algorithmic MTS and hardware GPGPU acceleration**, can significantly improve the overall performance of multiscale simulations.
- The performance improvements brought by MTS and GPGPU are both achieved through reducing the burden of force calculations on CPU thus they both suffer **communication bottleneck**.
- The rule of thumb is to consider the **balance of speed and accuracy** for an optimal MTS scheme and the **balance of computation and communication** for an optimal load-balancing scheme between accelerators and CPUs.
- Recommendations for future work: study more **complicated blood biological processes** and search for solutions of reducing communication bottleneck for the demands of **larger simulation system**.

Acknowledgments

Advisors and Fellows:

- Yuefan Deng, PhD (Department of Applied Mathematics and Statistics, Stony Brook University)
- Danny Bluestein, PhD (Department of Biomedical Engineering, Stony Brook University)
- Peng Zhang, PhD (Department of Biomedical Engineering, Stony Brook University)
- Jawaad Sheriff, PhD (Department of Biomedical Engineering, Stony Brook University)
- Seetha Pothapragada, PhD (Department of Applied Mathematics and Statistics, Stony Brook University)
- Chao Gao, M.S. (Department of Biomedical Engineering, Stony Brook University)
- Li Zhang, M.S. (Department of Applied Mathematics and Statistics, Stony Brook University)

Dissertation Committee:

- Professor James Glimm
- Professor Yuefan Deng
- Professor Robert Harrison
- Professor Danny Bluestein

Computing Resources:

- Seawulf Cluster (Stony Brook University)
- Sunway Blue Light System (National Supercomputing Center in Jinan, China)
- Stampede System (Texas Advanced Computing Center, under XSEDE project)
- Tianhe-2 (National Supercompuing Center in Guangzhou, China)
- Cray CS-Storm (Cray Inc.)

Support Funding:

- NIH (NHLBI R21 HL096930-01, NIBIB Quantum U01EB012487, DB)

Conference Travel Funding:

- Institute for Advanced Computational Science (IACS)
- Department of Applied Mathematics and Statistics, Stony Brook University