Mapping Biomedical Applications onto GPU Platforms

Joseph JaJa
University of Maryland

Fluid-Structure Interactions

- Collaboration between GWU (Balaras), UMD (Solares, Wu), and University of Chicago (Dubey).
- Goal: Development of high performance algorithms applicable to fluid-structure interactions in viscous incompressible flows.
- Example application: interactions between the red blood cells and plasma
- Critical Components: Poisson equation solver combined with a multigrid algorithm. Multi-dimensional FFTs and several types of matrix computations

Data-Driven Understanding of Brain Disorders

- Collaboration between the University of Maryland (Varshney and JaJa) and the University of Maryland at Baltimore (Gullapalli, Herskovits, etc.)
- Understanding of brain connectivity differences between subjects with brain disorders and normal subjects using diffusion MRI.
- Dynamics of functional brain connectivity using resting state fMRI, for subjects with moderate TBI.

Connectivity Matrix

- Diffusion MRI images with 64 diffusion frames with resolution 128×128×52.
- Probabilistic Tractography
- Number of entries in the sparse connectivity matrix: 100,000,000-200,000,000.
- Number of voxels in ROI: 100,000-200,000.

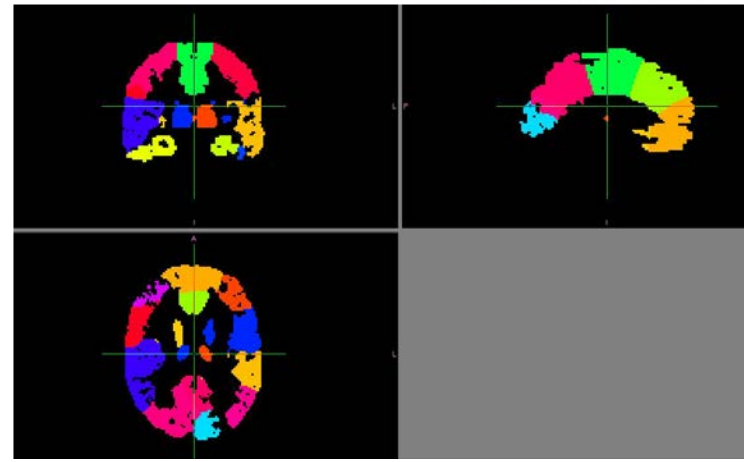

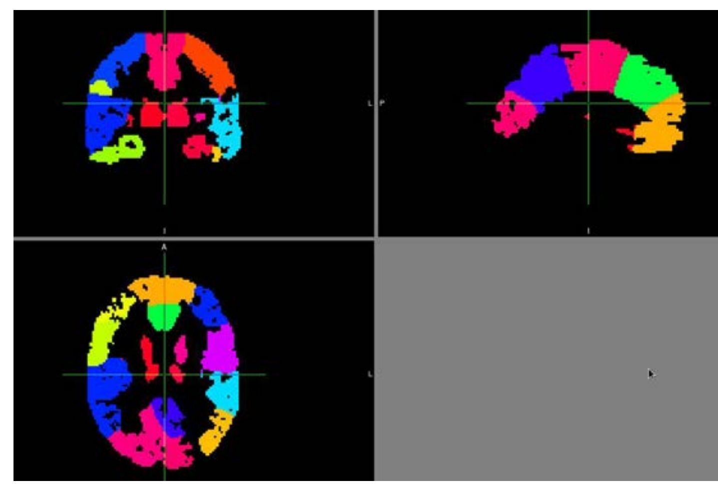

Inflammatory Responses and Wound Healing in Vocal Fold

N. Seekhao
Collaborators: N. Li, C. Shung, L. Mongeau (McGill U.)

Biomechanical Stress

↓

Mucosal Damage

↓

Cell Recruitment

↓

Cell Function

Inflammatory Responses & Wound Healing in Vocal Fold

Biomechanical Stress

↓

Mucosal Damage

↓

Cell Recruitment

↓

Cell Function

Force applied on tissue. Talking, shouting etc.

Image from : http://2.bp.blogspot.com/-DI0yRAeRKjA/TrDdREMzr_I/AAAAAAAAH6o/QFgZ7xFFRjg/s320/shout.png

Inflammatory Responses & Wound Healing in Vocal Fold

Biomechanical Stress

↓

Mucosal Damage

↓

Cell Recruitment

↓

Cell Function

Damage in the tissue of the vocal fold

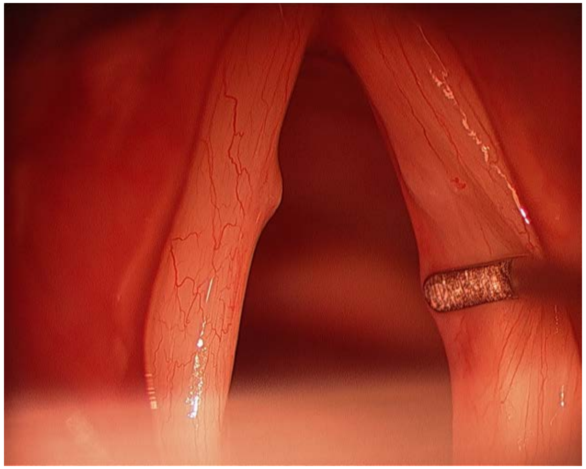

Image from :
https://wiki.uiowa.edu/download/attachments/39001206/nodules%20op%205.png?api=v2

Inflammatory Responses & Wound Healing in Vocal Fold

Biomechanical Stress

↓

Mucosal Damage

↓

Cell Recruitment

↓

Cell Function

Attracting cells such as platelets, neutrophils, and macrophages to the wound site

Image from :
http://www.biospectrumasia.com/IMG/362/44362/atherosclerotic-lesions-generates-robust-t-cell-anti-inflammatory-response-262x174.jpg

Inflammatory Responses & Wound Healing in Vocal Fold

Biomechanical Stress
↓
Mucosal Damage
↓
Cell Recruitment
↓
Cell Function

Each cell perform its duty. One or more of the following:

- Secrete chemical (IL-1, MMP-8 etc.) to attract, excite or inhibit other cells
- Deposit ECM protein (collagen, elastin etc.) to heal damaged tissue
- Clean up cell debris

Agent-Based Modeling (ABM)

1. Bottom-up, rule-based, discrete-event and discrete-time computational model
2. Initial "World" and a collection of "agents."
3. Interactions between agents and the world.
   - Agents migrate to area of injury
   - Remove dead cells and tissue debris
   - Remodel ECM to heal damaged tissue
4. Stochastic moves
5. Emergent behavior

ABMs of Vocal Fold Wound Healing Process

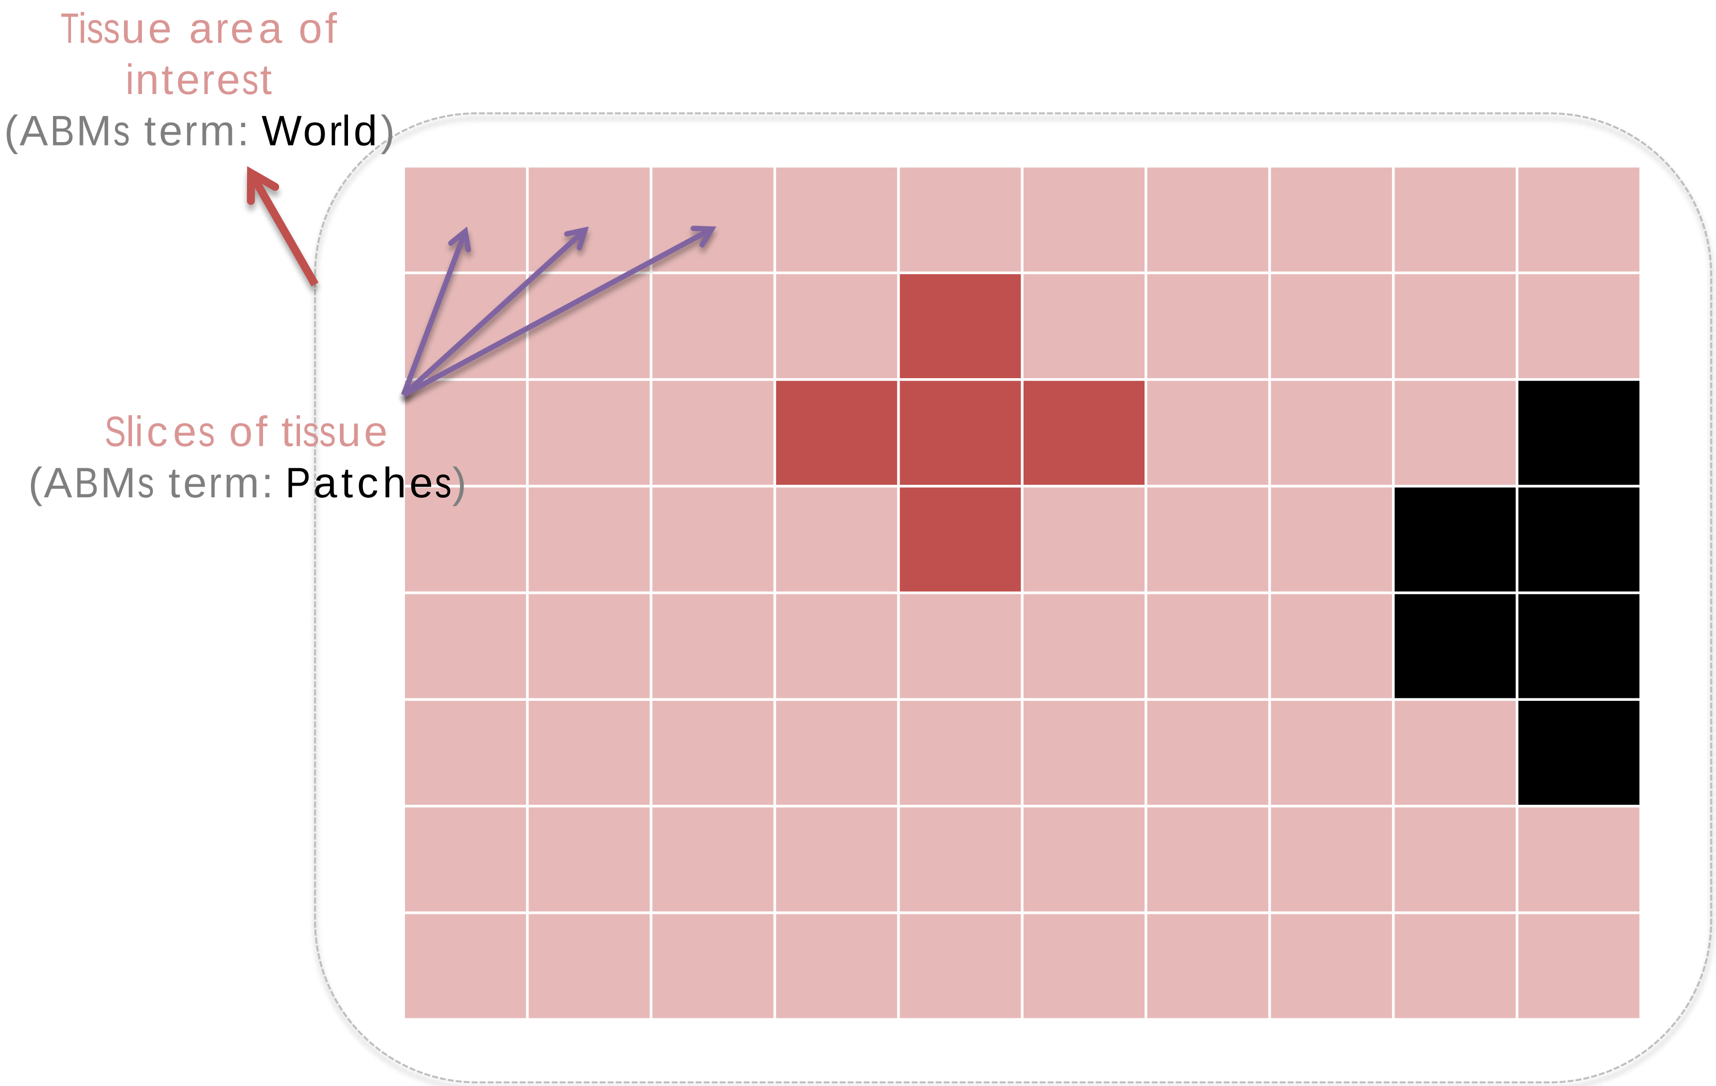

ABMs of Vocal Fold Wound Healing Process

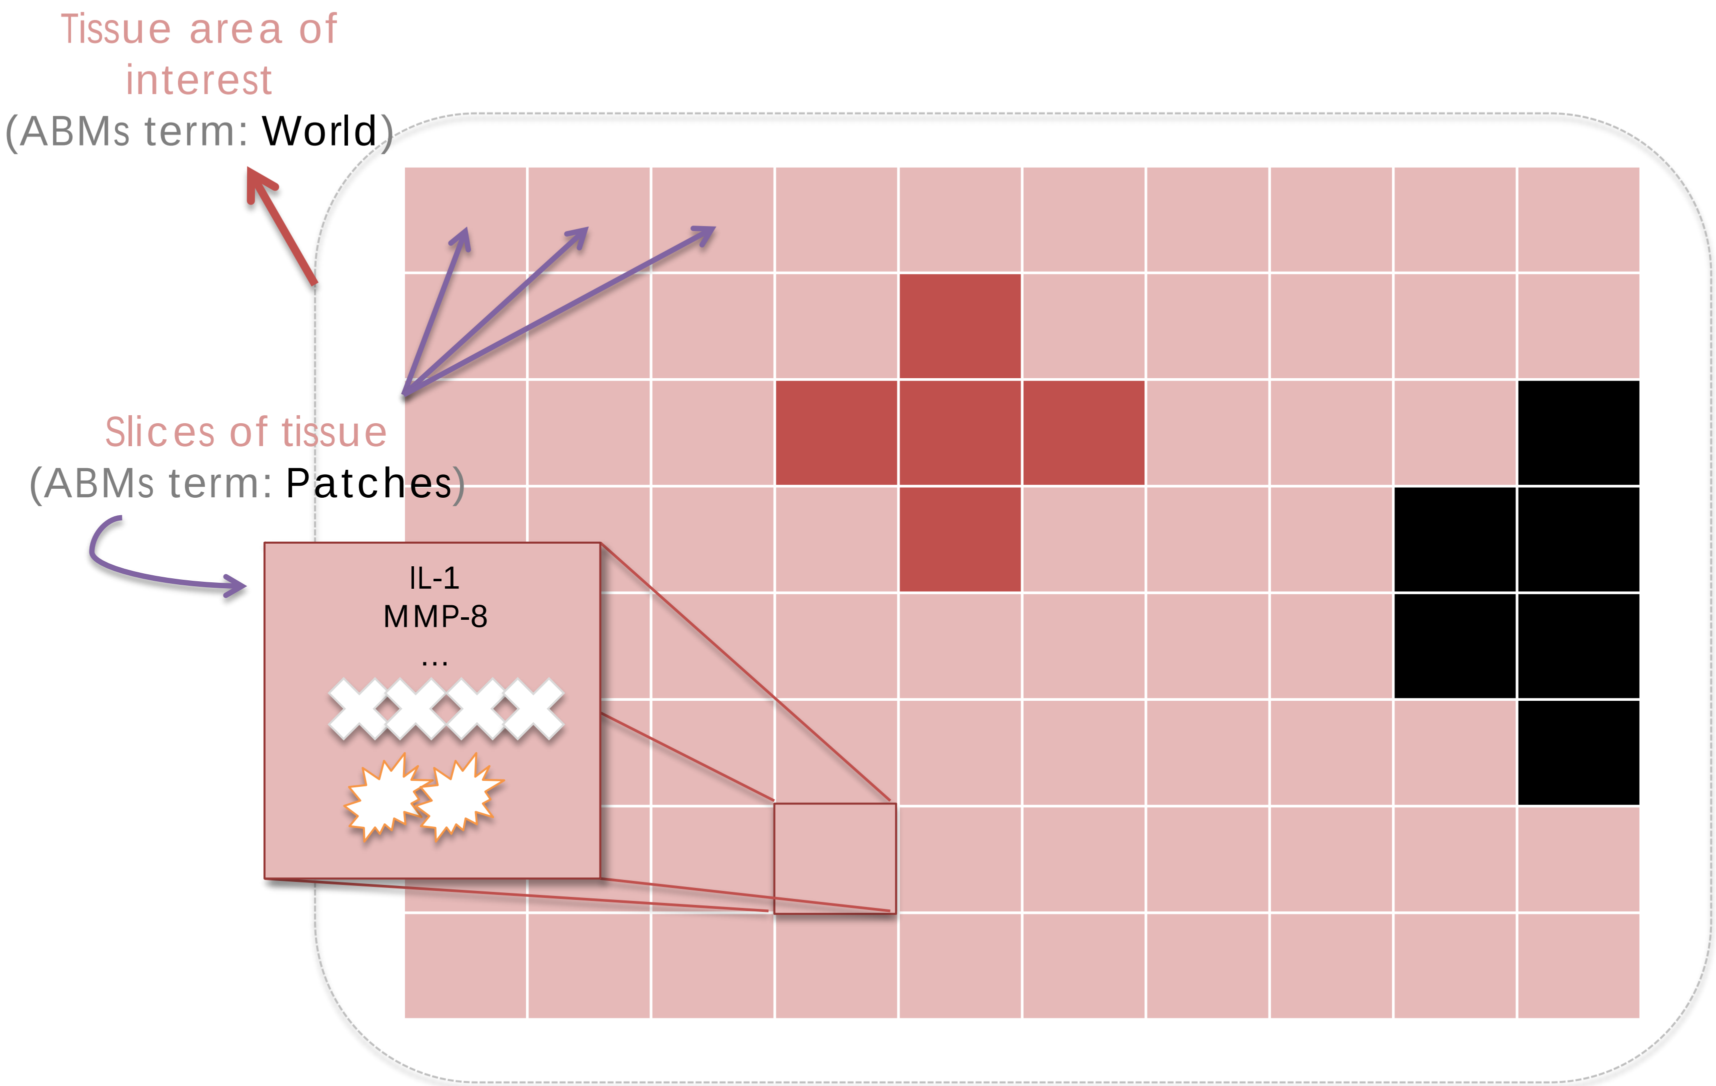

ABMs of Vocal Fold Wound Healing Process

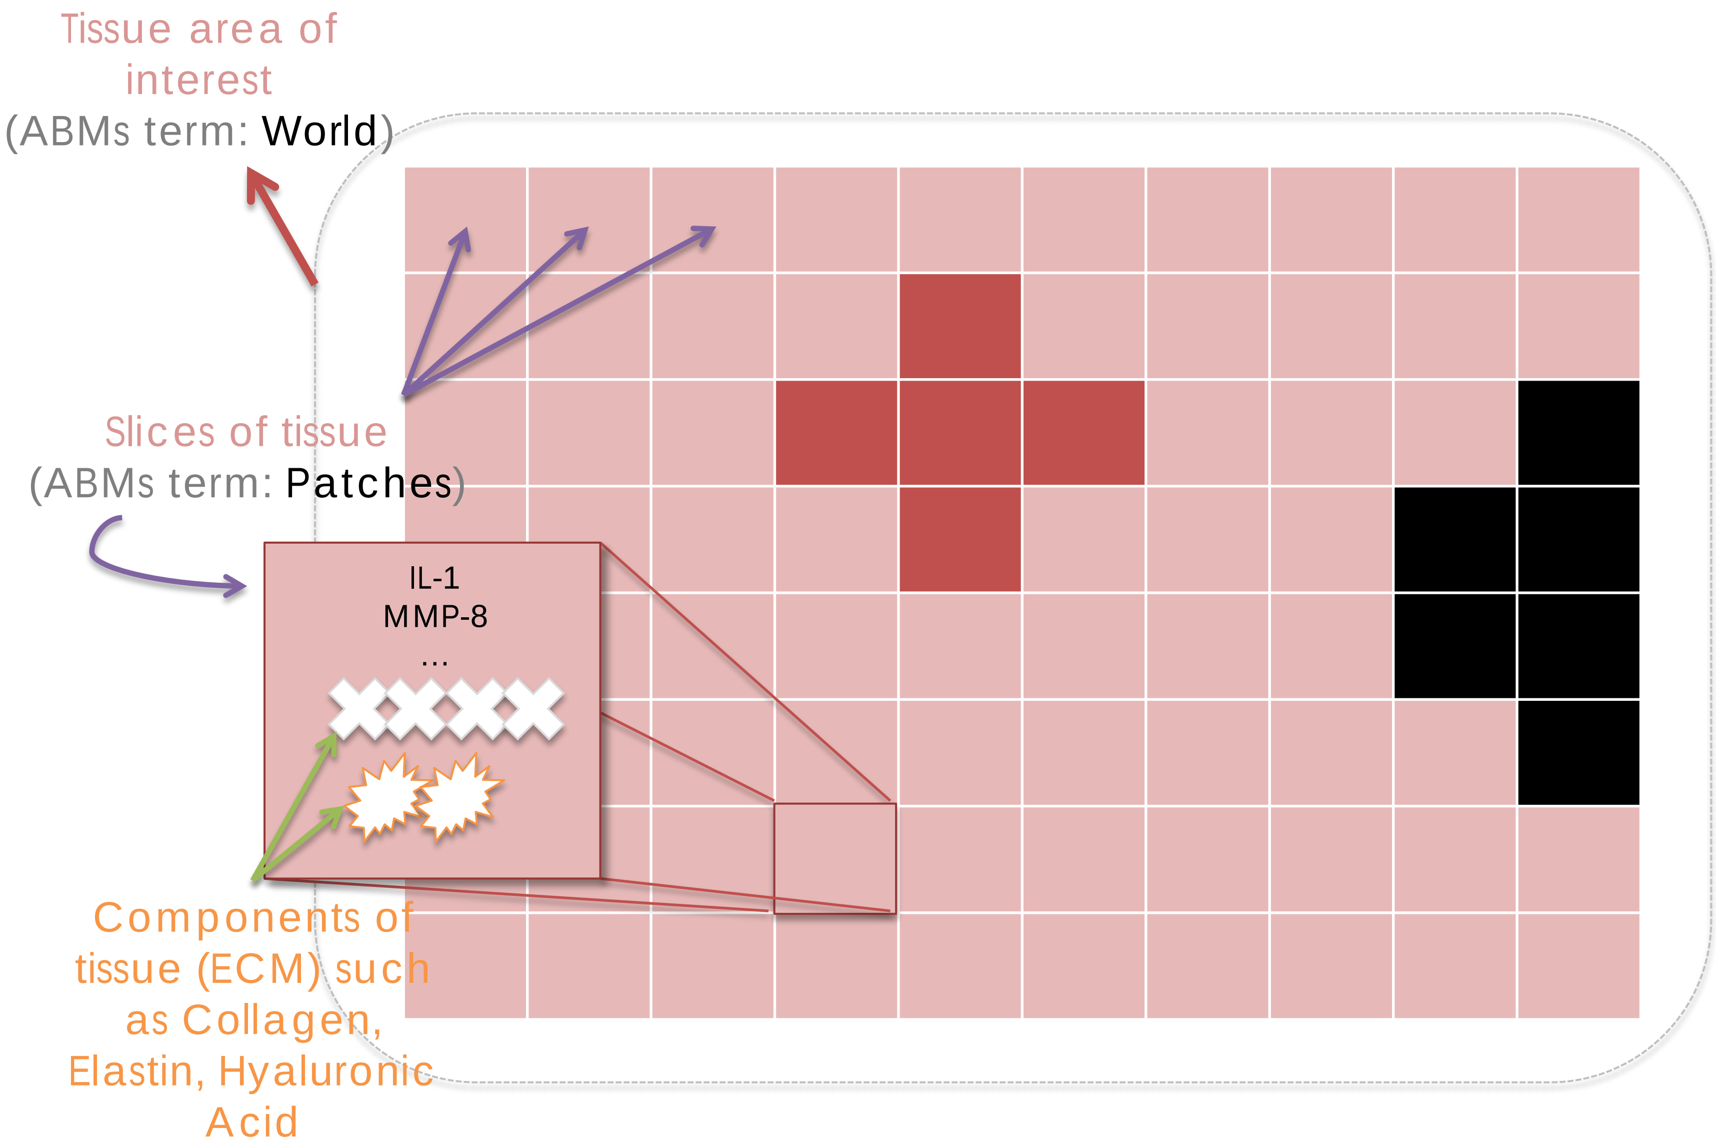

ABMs of Vocal Fold Wound Healing Process

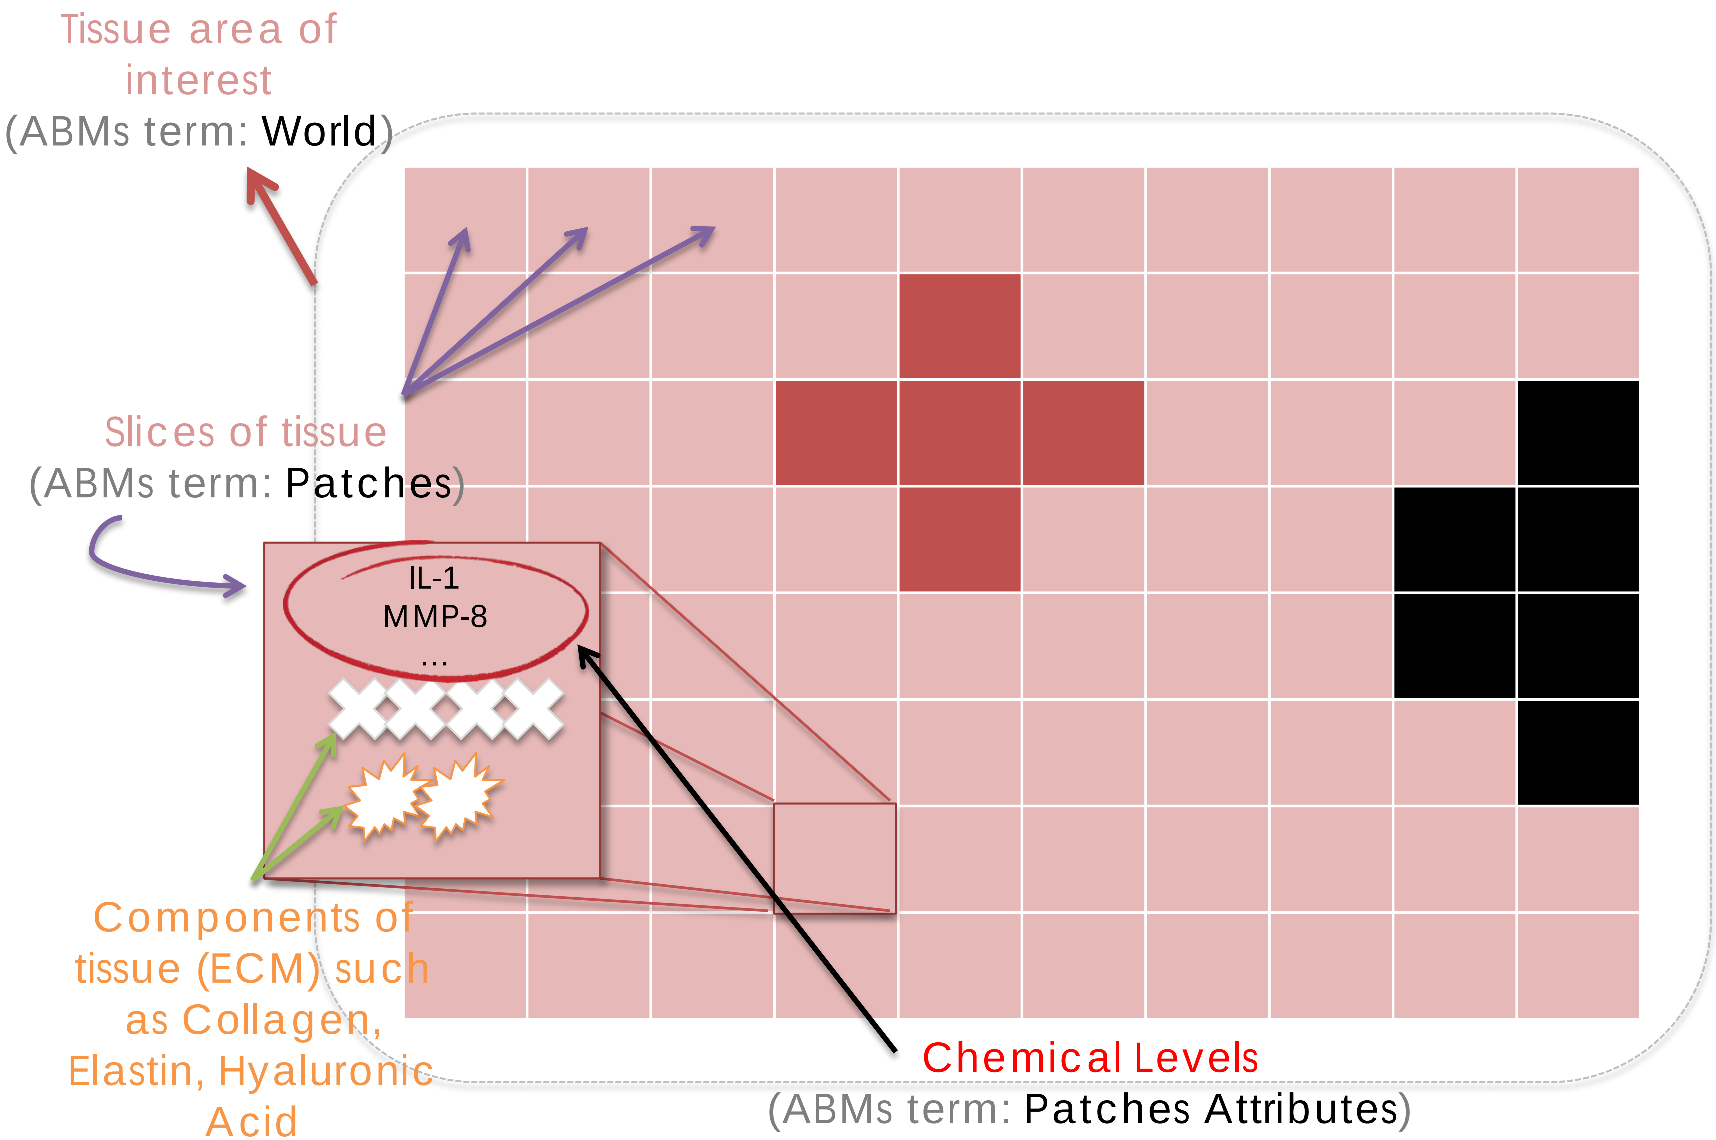

ABMs of Vocal Fold Wound Healing Process

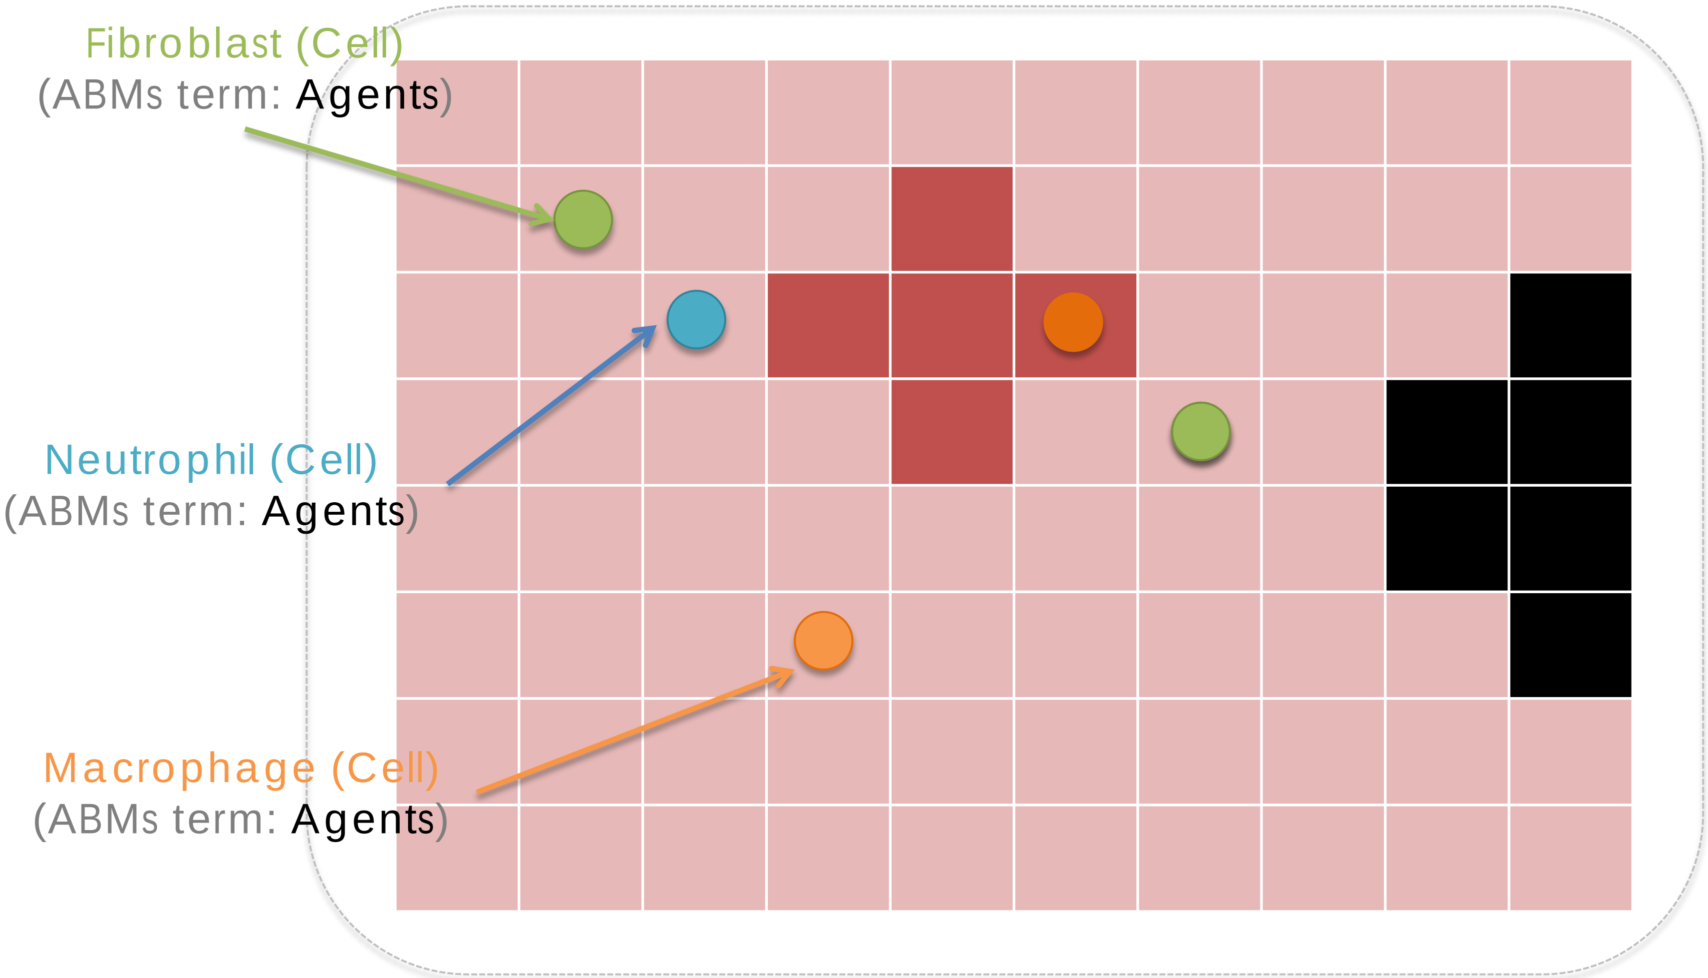

ABMs of Vocal Fold Wound Healing Process

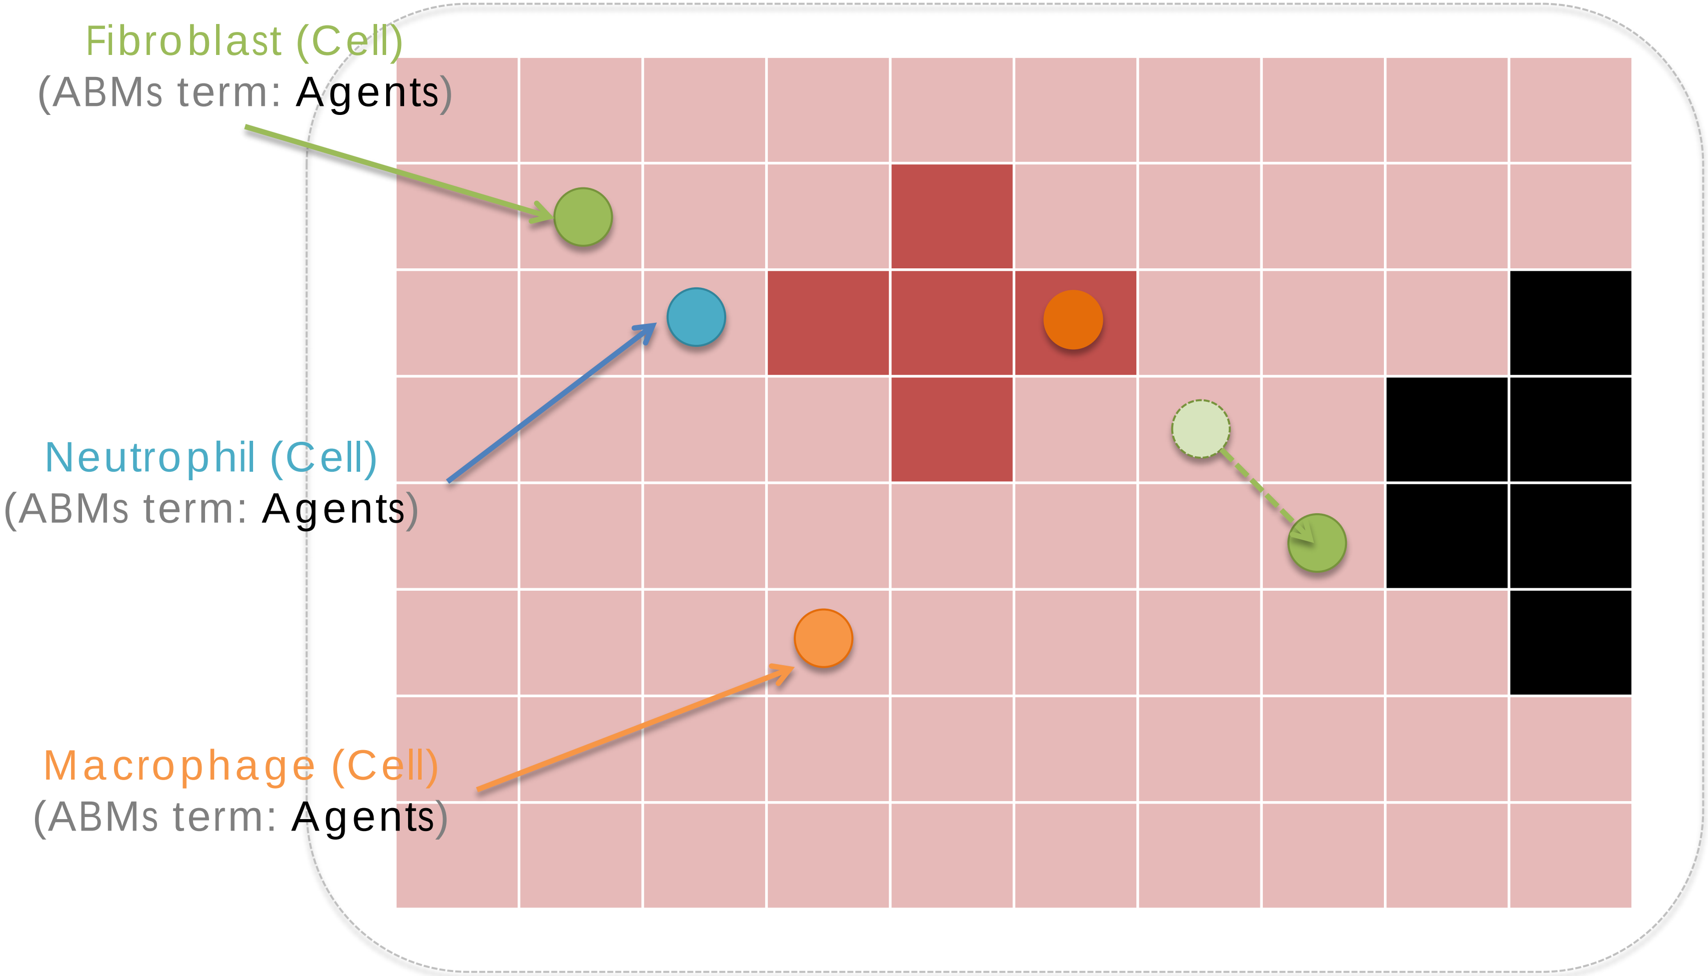

ABMs of Vocal Fold Wound Healing Process

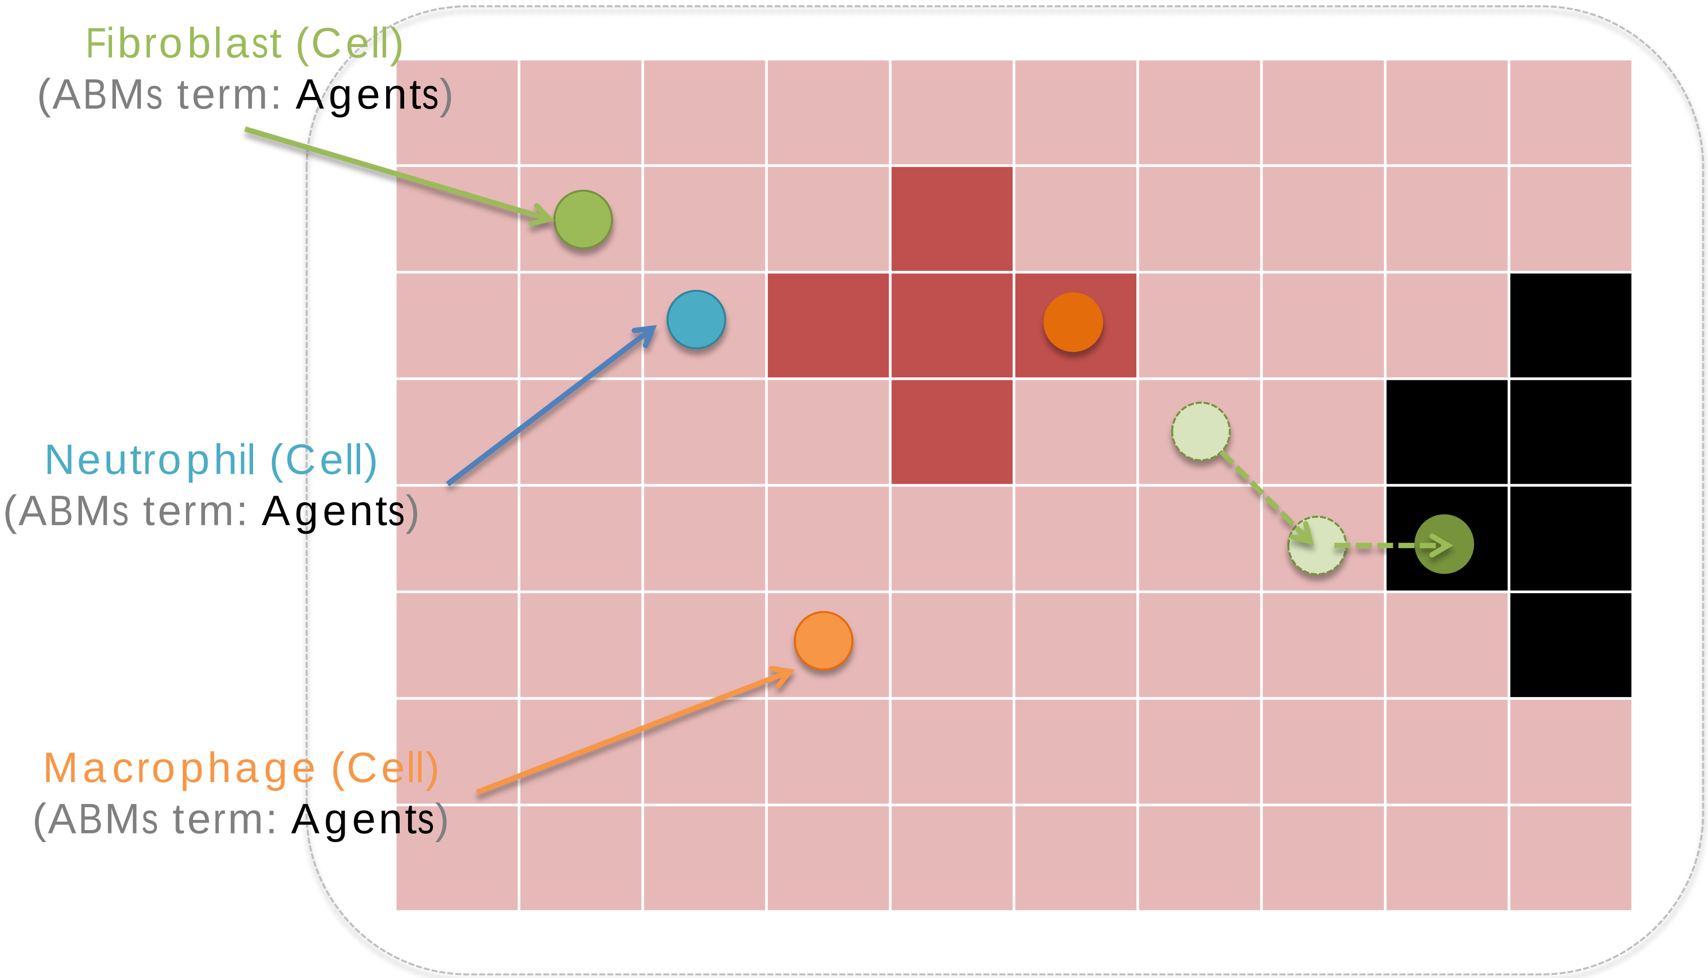

Problem Scale

	Unit	Current Model	New Model	3D
World size	patches x patches	120 x 240	1160 x 1660	1160 x 1660 x 100
Patch size	mm x mm	0.04 x 0.04	0.015 x 0.015	0.015 x 0.015 x 0.015
Simulated area dimension	mm x mm	4.8 x 9.6	17.4 x 24.9	17.4 x 24.9 x1.5
Number of Patches		28.8k	1.9m	192.56m
Simulated area size	mm^2	46.08	433.26	n/a
Simulated area volume	mm^3	n/a	n/a	649.89
Initial number of:				
Neutrophils	cells	100	182.4k	18.24m
Macrophages	cells	50	22.8k	2.28m
Fibroblasts	cells	50	22.8k	2.28m

* Number of cells increase throughout the simulation due to proliferation. Current model shows a doubling of number of cells after the end of “5-day” simulation.

Characteristic Features of Applications

- Computationally demanding applications with irregular memory access patterns and involving large data sizes that cannot fit on the GPU memory
- Need to use heterogeneous platforms involving multicore CPU with one or more many-core GPUs.
- Performance Goal: try to achieve the same performance rate or throughput as in the case when the data fits on the GPU.

Heterogeneous Platforms

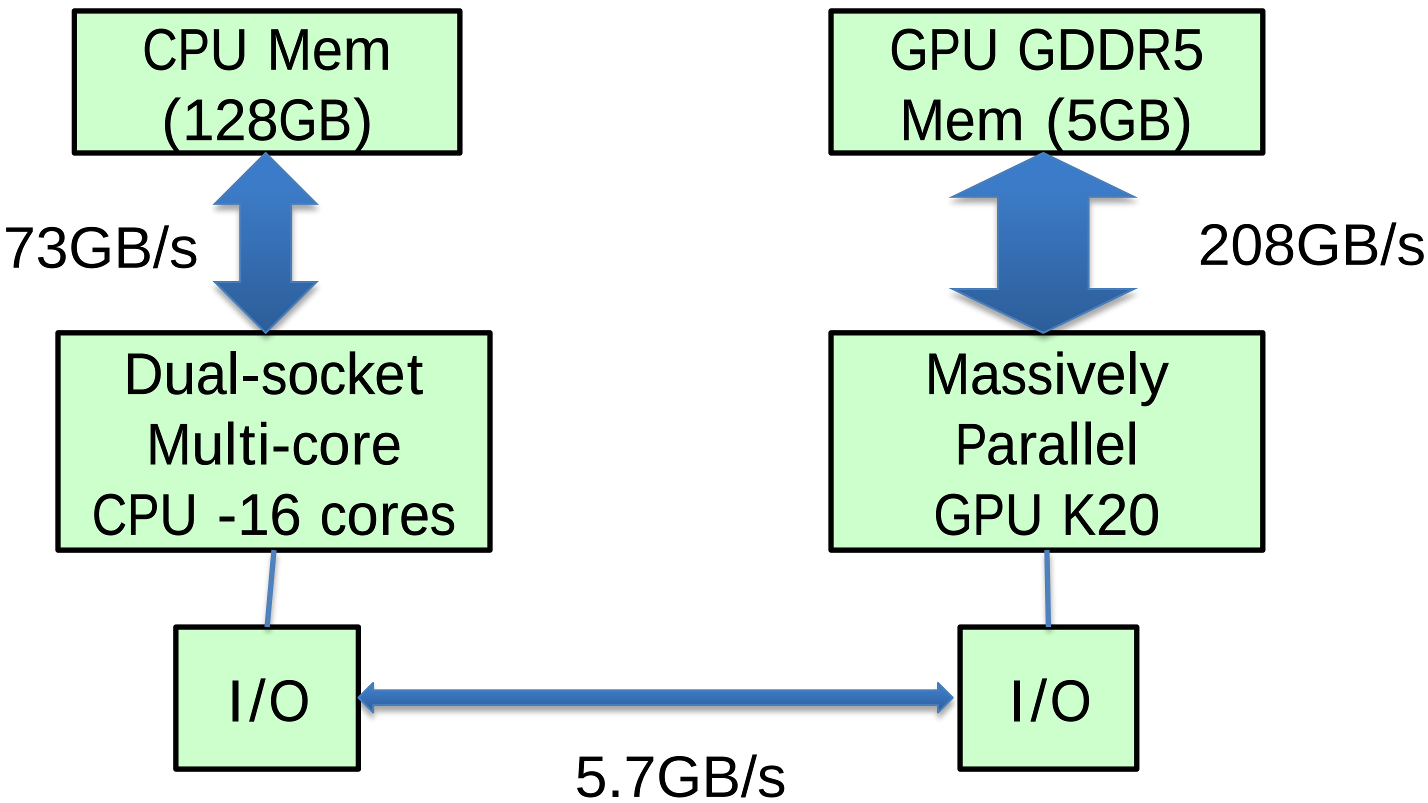

Dense Matrix Multiplication

- DGEMM: $$C = \alpha AB + \beta C$$
  where the matrices are of dimensions:
  $$A : m \times k, B : k \times n, C : m \times n$$
- Why?
  - An important kernel for many problems
  - Optimization ideas can be used in other problems
  - Perhaps the most-studied algorithm in high performance computing
- Can we solve very large DGEMM with the same performance throughput as small DGEMM?

Block Matrix Multiplication

- Decompose into blocks

$$C_{ij} = \alpha \sum_{k=0}^{K/bk} A_{ik}^{(0)} B_{kj}^{(0)} + \beta C_{ij}^{(0)}$$

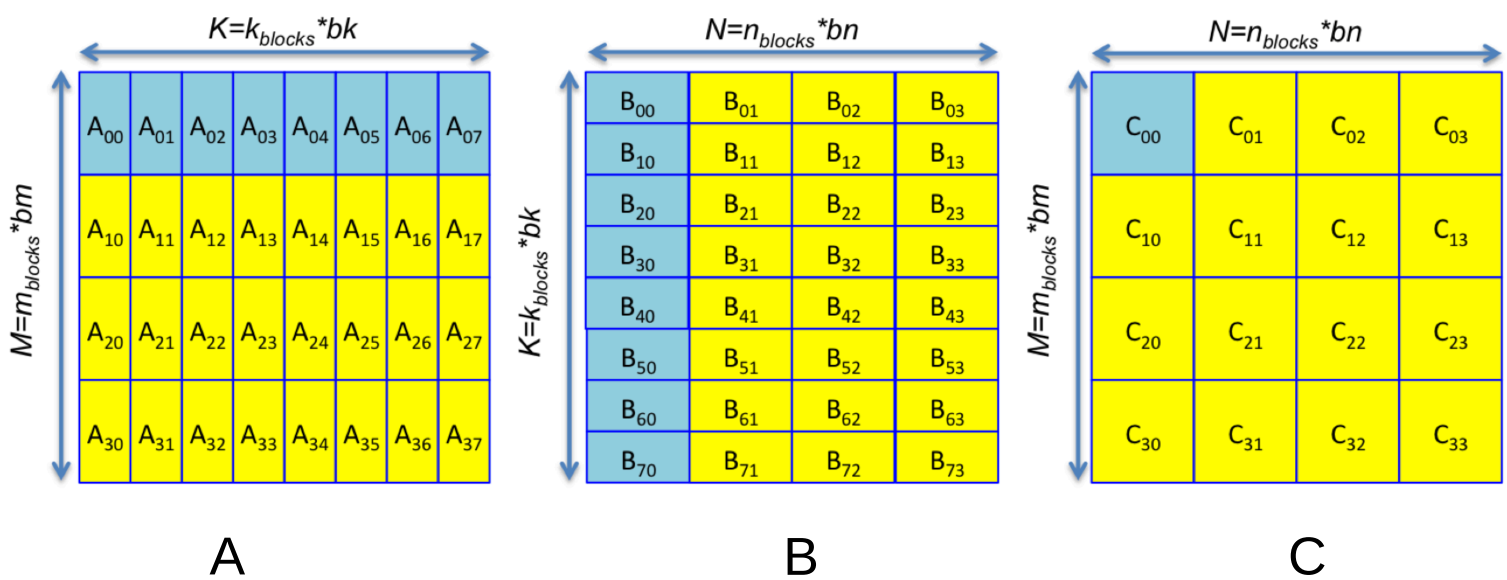

Multiple CUDA Stream Scheduling

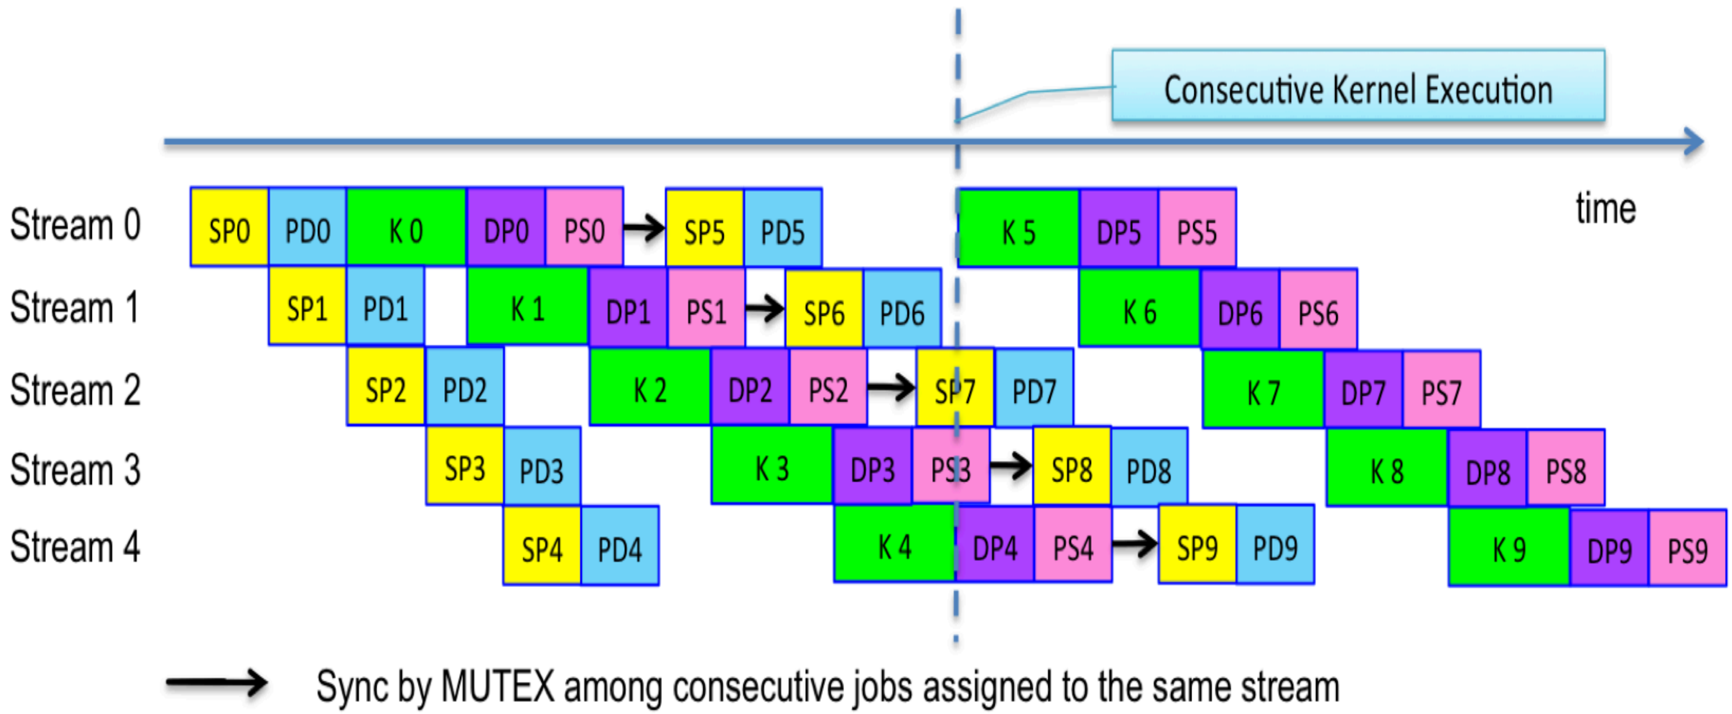

Figure 6: CPU-GPU Software Pipeline

Performance Evaluation

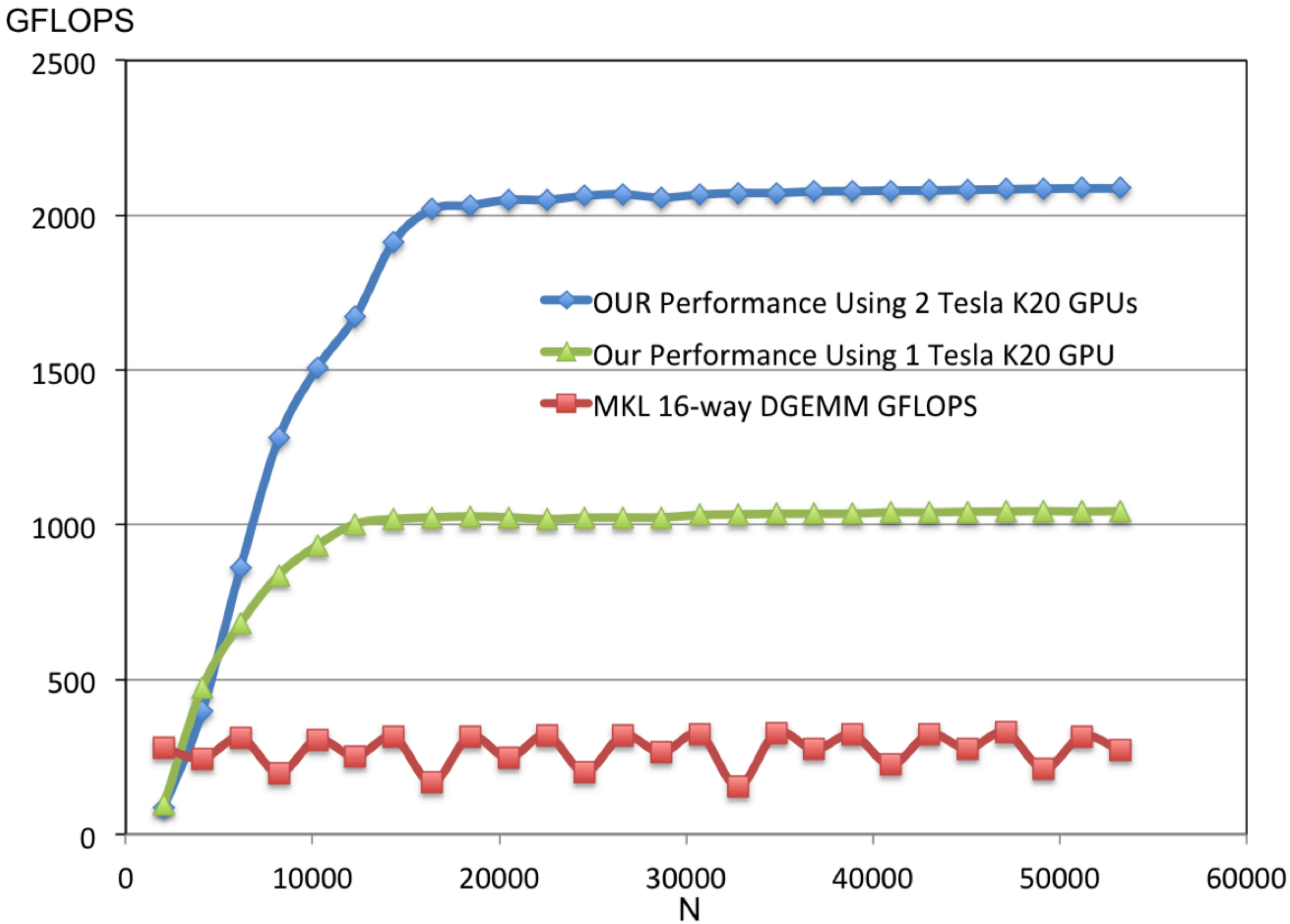

Figure 7: DGEMM Performance on *Sandy-Kepler* Node

Performance Evaluation

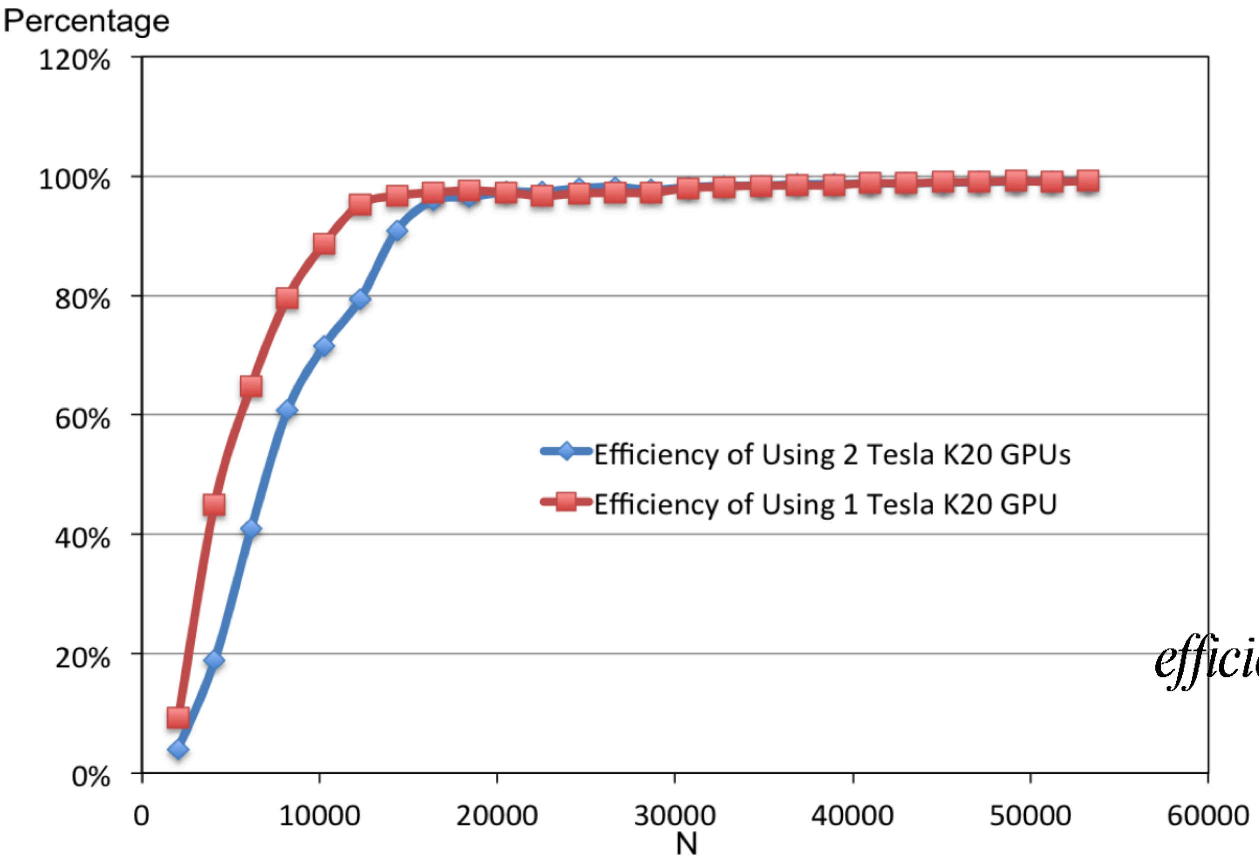

$$efficiency = \frac{GFLOPS_{achieved}}{GFLOPS_{peak\,lib\,performance}}$$

Figure 9: Efficiency on *Sandy-Kepler* Node

Conclusion

- Many biomedical applications can make effective use of heterogeneous platforms.
- But a significant amount of work is required to organize the computation into multi-stream of data transfers and kernel executions with no or very small stall time.
- Portability of high performance code remains a problem.