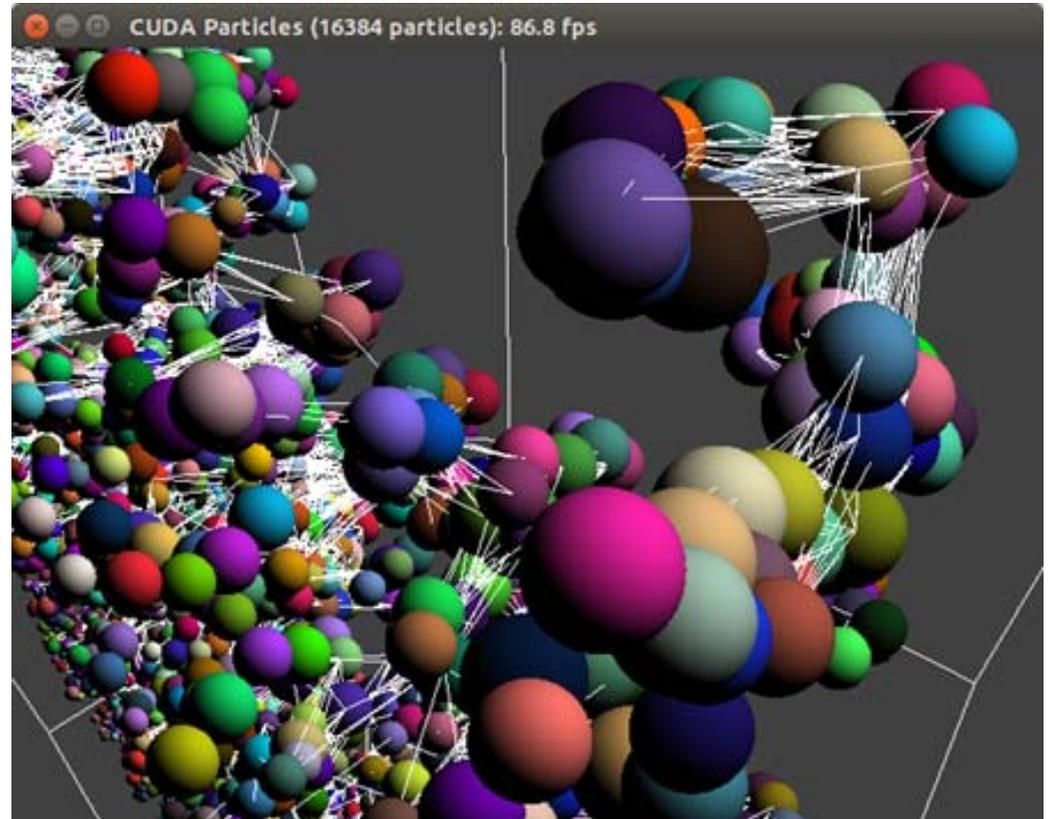


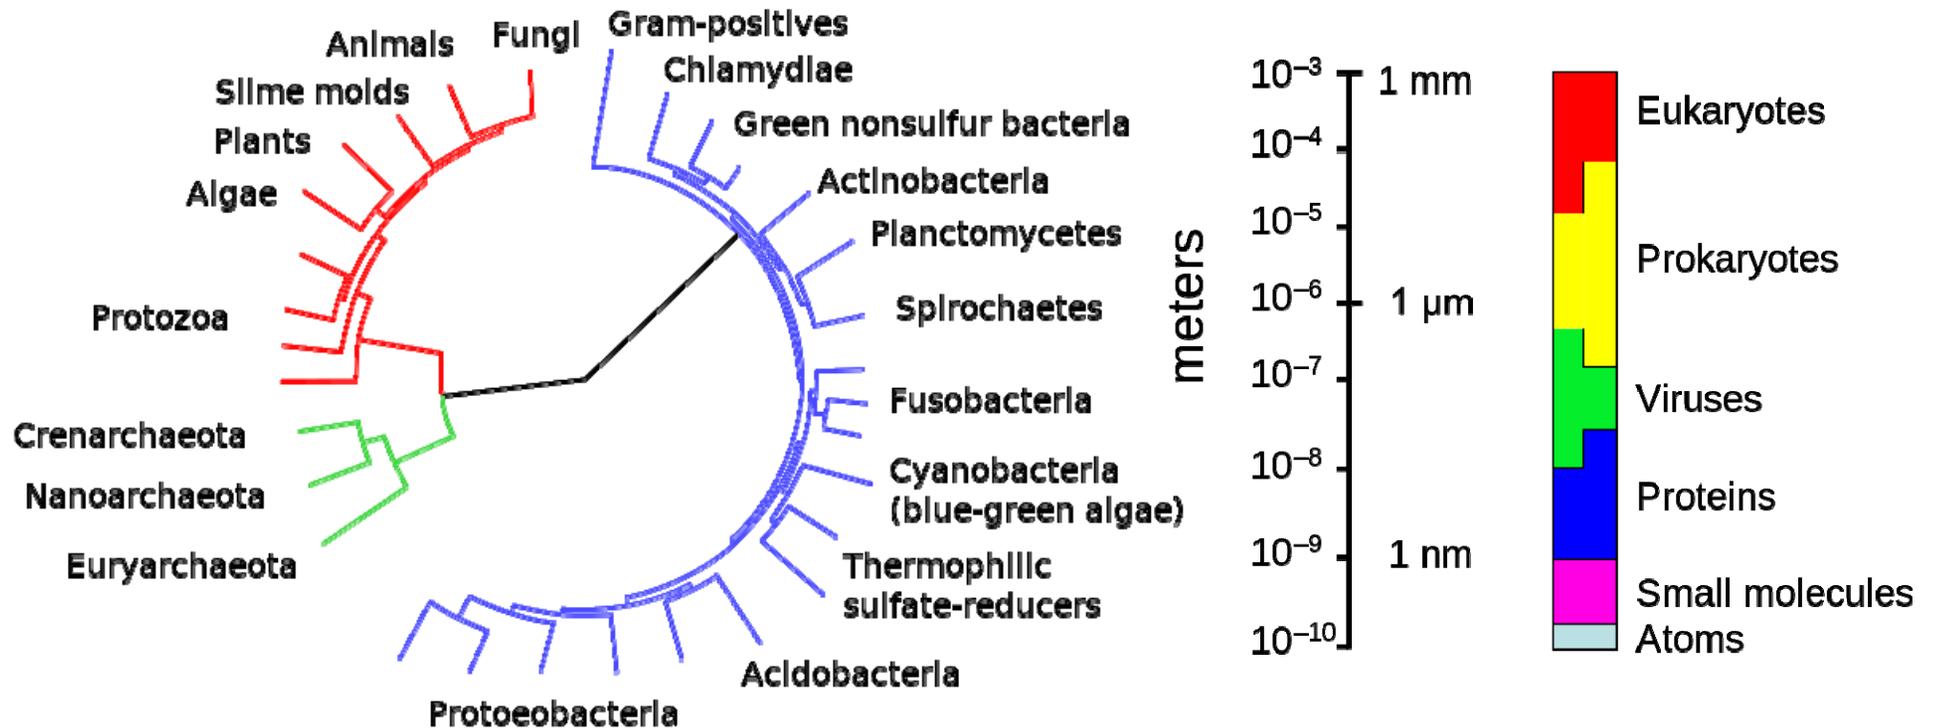
Fighting Infections and Antimicrobial Resistance Through GPU-Accelerated In Silico Models

Radu Marculescu
Carnegie Mellon University

GTC, April 5, 2016

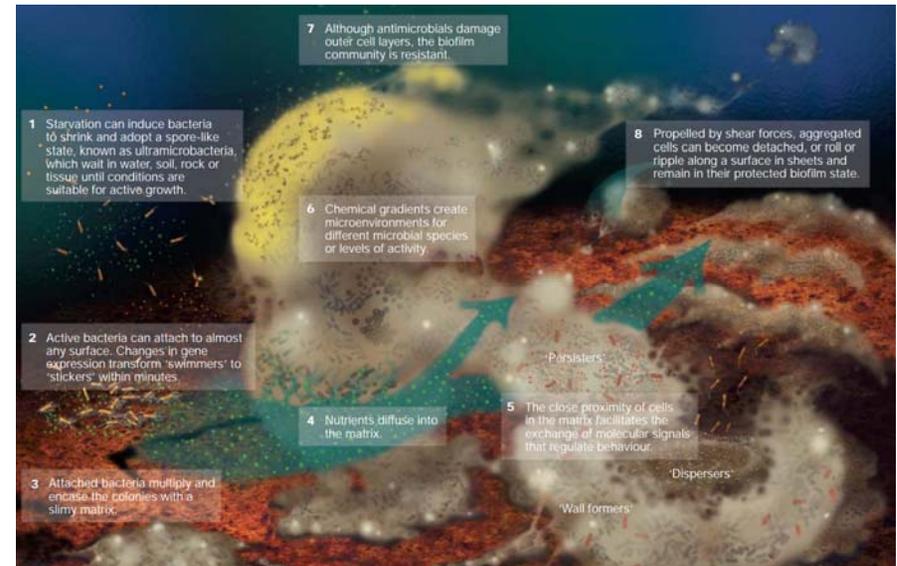
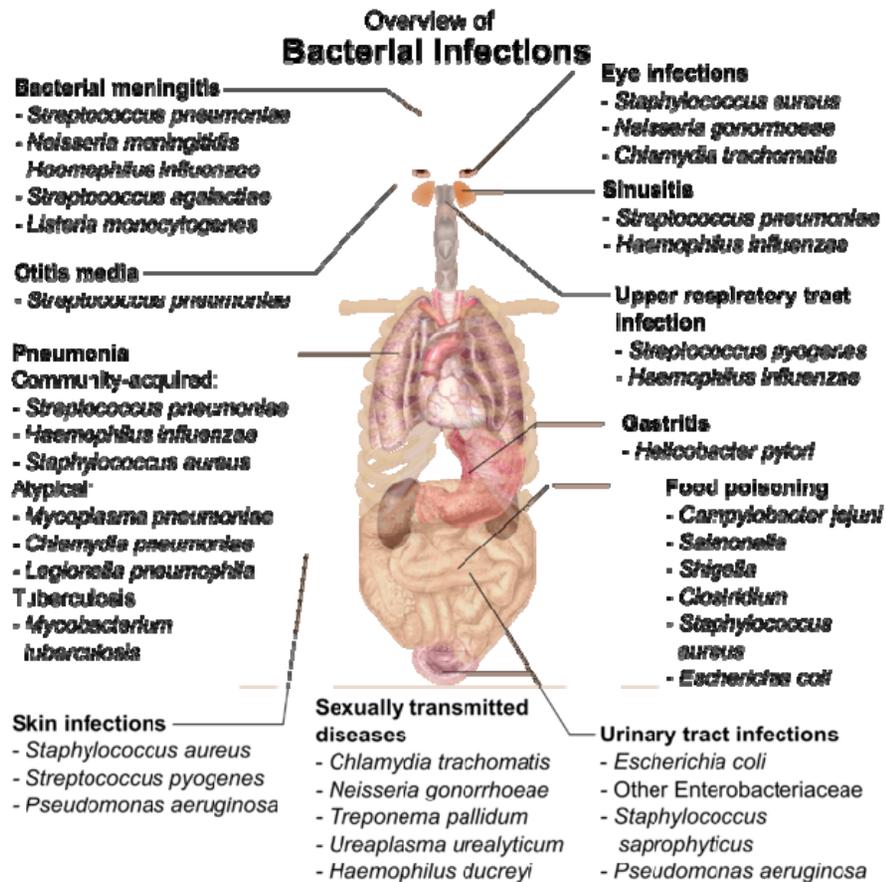


Bacteria are among the earliest forms of life that appeared on Earth billions of years ago. Studied since 1676...



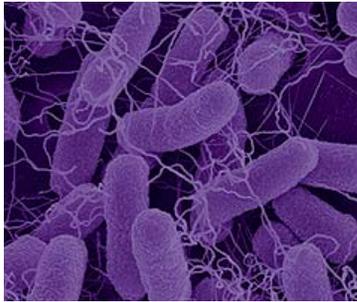
Bacteria are of immense importance because of their extreme flexibility, capacity for rapid growth and reproduction.

The fight against bacterial infection represents one of the big challenges of modern medicine.

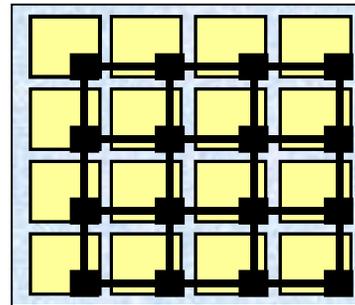


Focus on intercellular communication to understand behavior at population-level

This presentation focuses on the “network” paradigm and ways to engineer communication



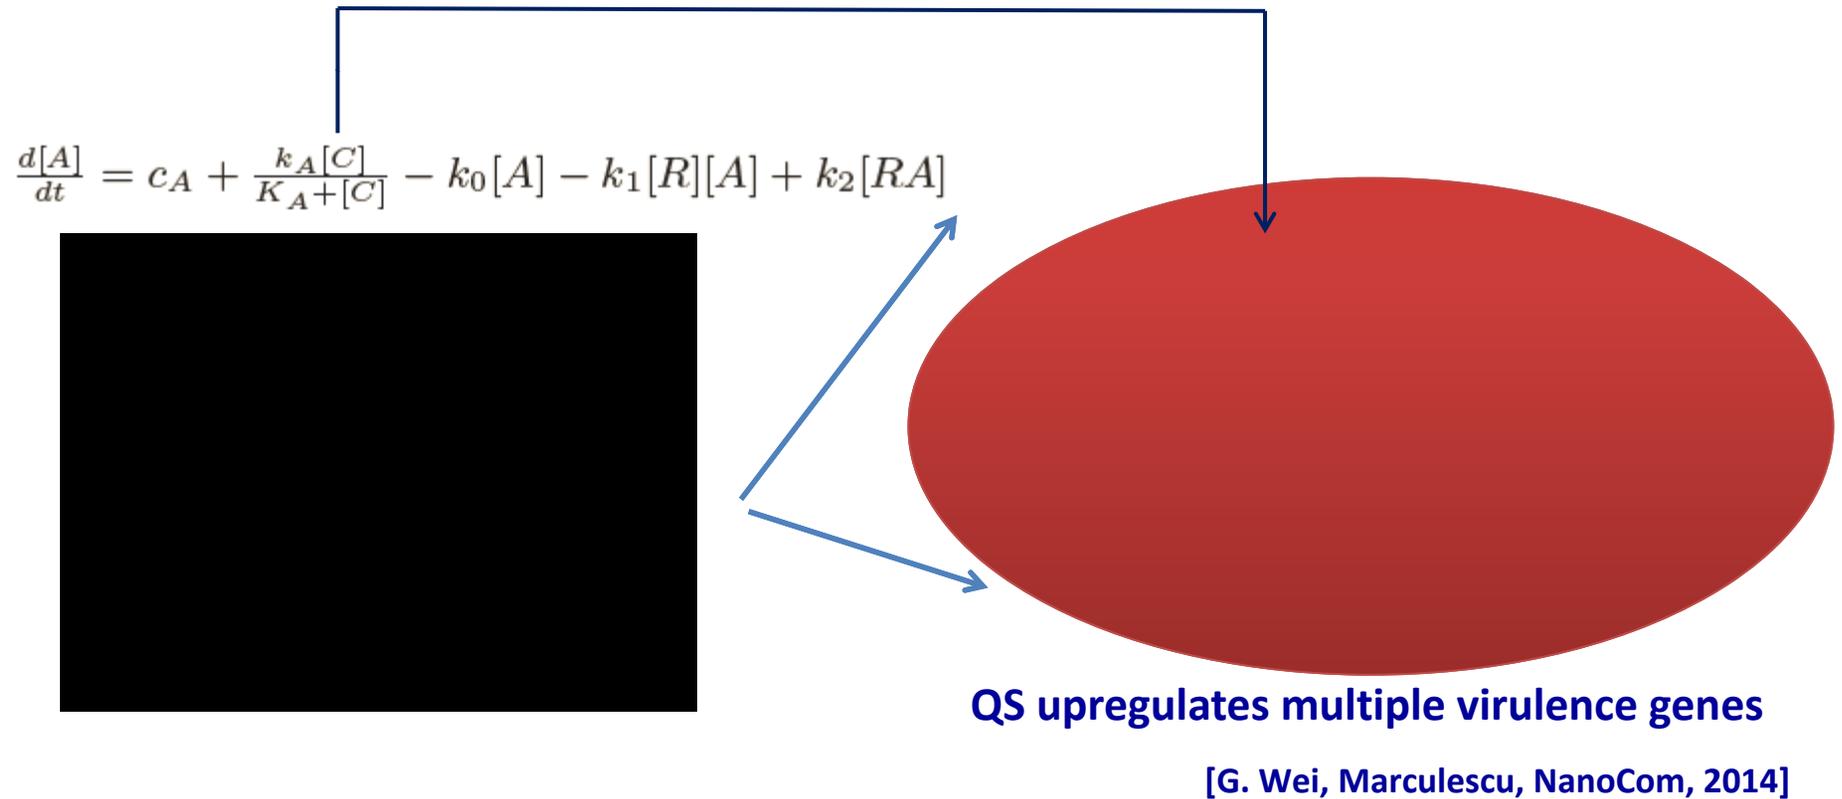
Biological networks
Cellular communication



Silicon networks
On-chip communication

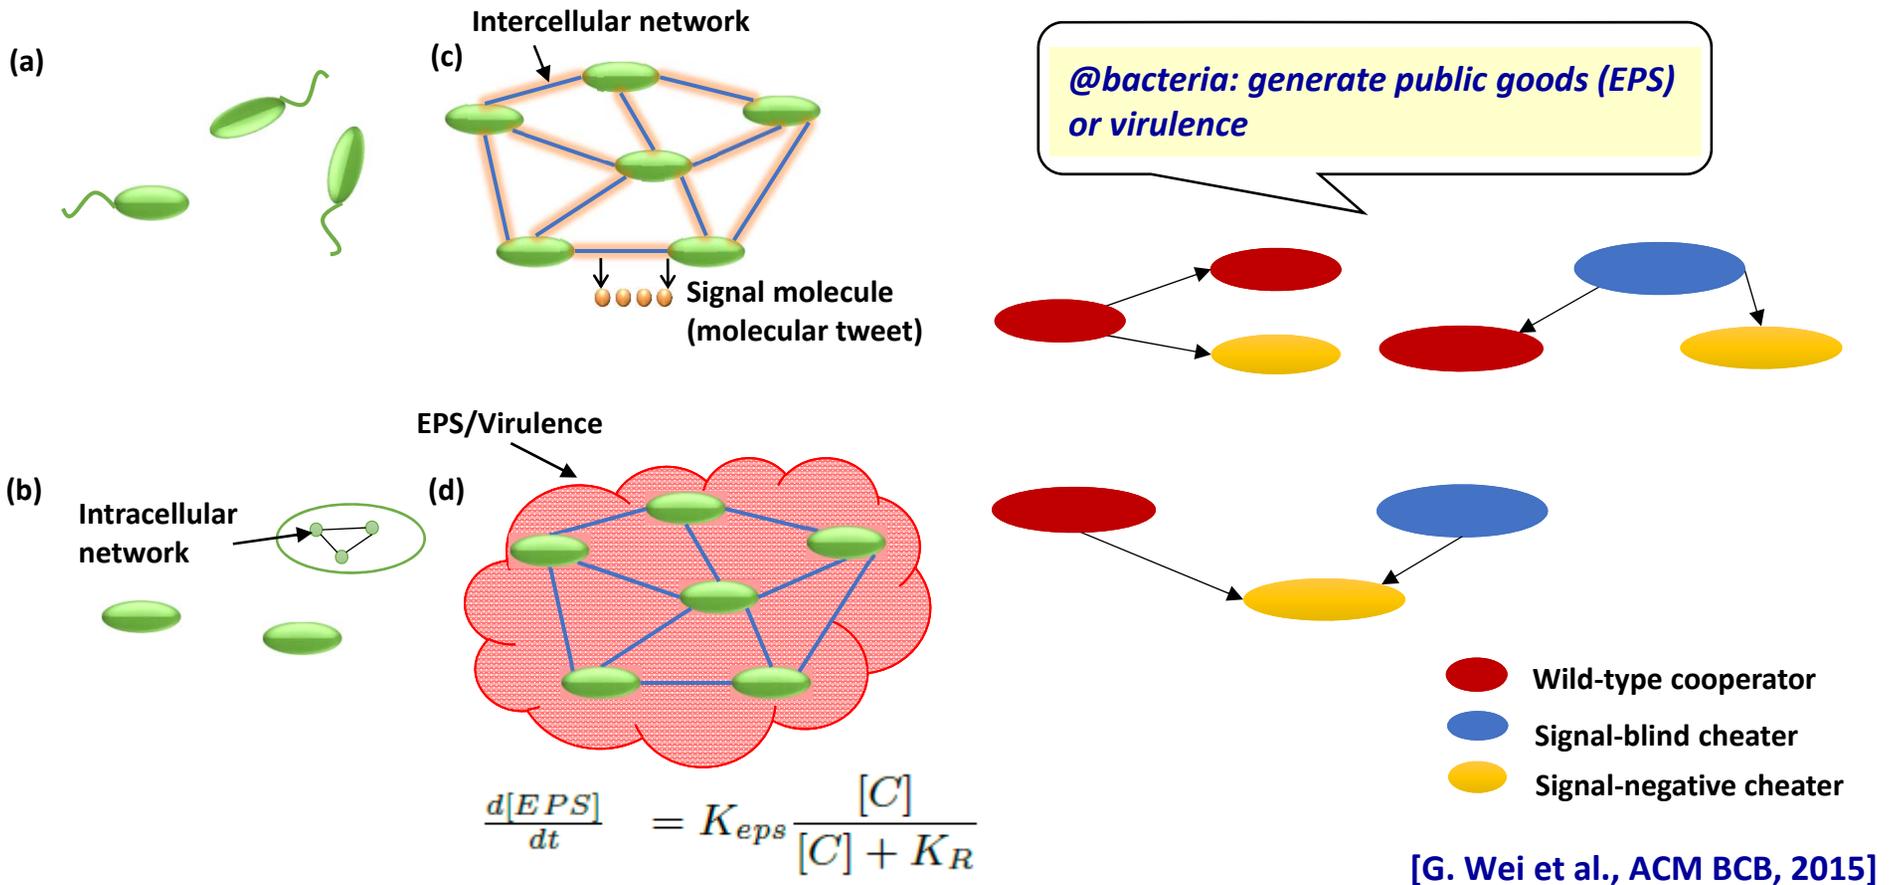
6 Bacteria form a “social network” to coordinate their behavior, build the biofilm and resist treatment.

A New Frontier – Quorum Sensing. QS is a density-dependent collective behavior that enables communication



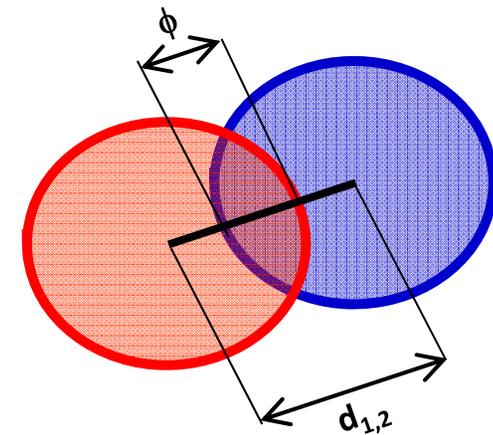
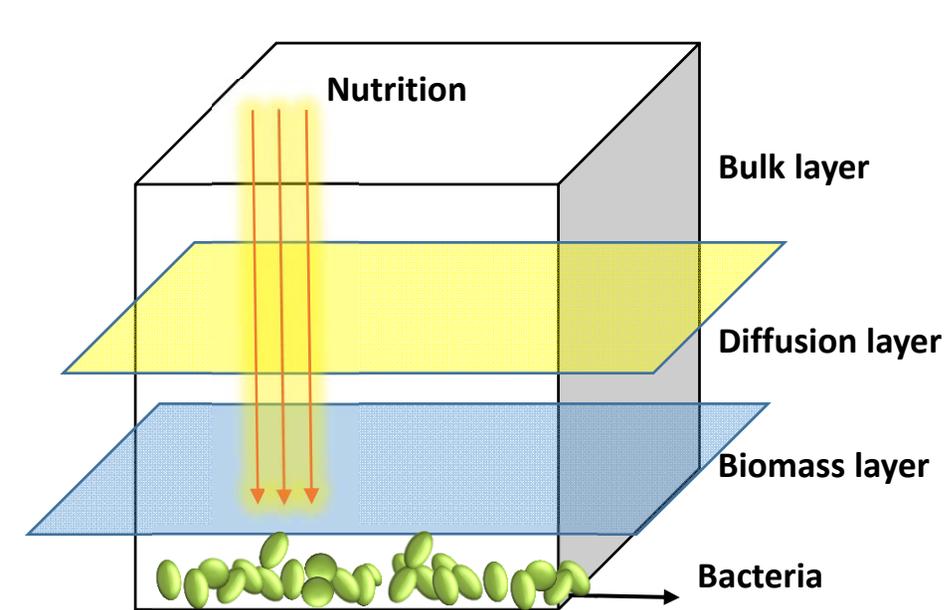
7 Gram-negative bacteria use largely homologous QS networks, where the AIs are detected and regulated via genetic circuits.

Bacterial biofilm: Use Twitter-like metaphor to explain participants and network formation

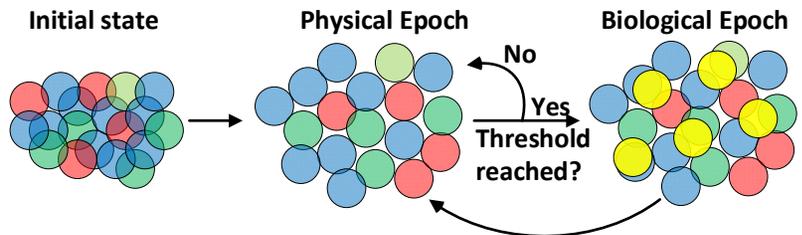
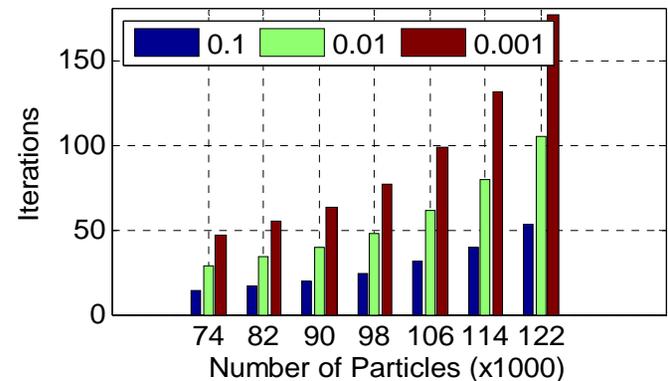


The structure and dynamics of the inter-cellular communication network is heavily influenced by its environment

Bacterial population dynamics is a complex problem. Direct wet-lab experimentation is costly and often impractical.

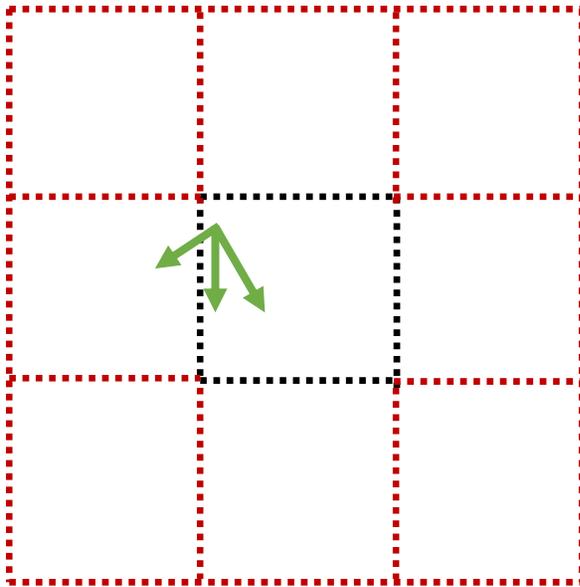


High-Accuracy = More iterations

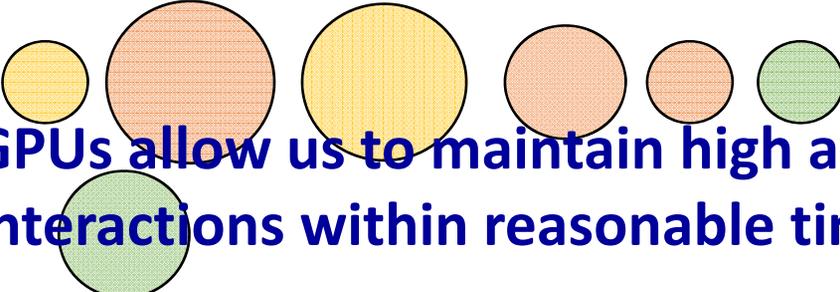


Experimenting with efficient GPU kernels using NVIDIA Thrust allows to achieve 100x acceleration with GTX980 GPUs.

Parallelize all agents on the GPU. Execute many of these physical interactions concurrently



- Arrange the data by physical location
 - `thrust::sort_by_key`
- For each **agent**
 - Examine **adjacent grid space**
 - Utilize `__shared__` cuda memory
 - Determine **movement vectors** between **agent** and **all other cells**
 - Move agent with movement vector
- Repeat until system approaches steady state
 - `thrust::transform_reduce` to get total movement in the system

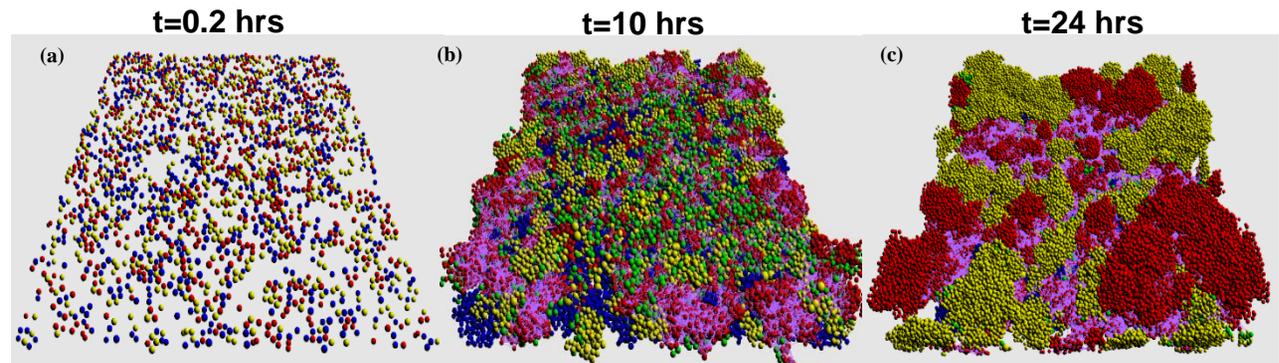
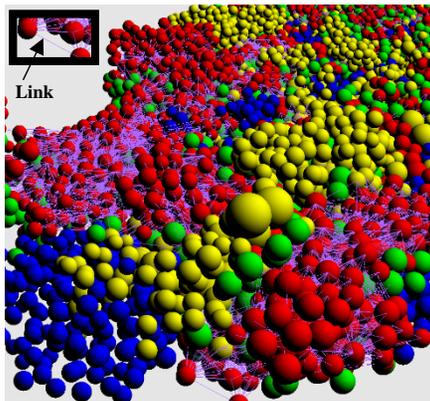
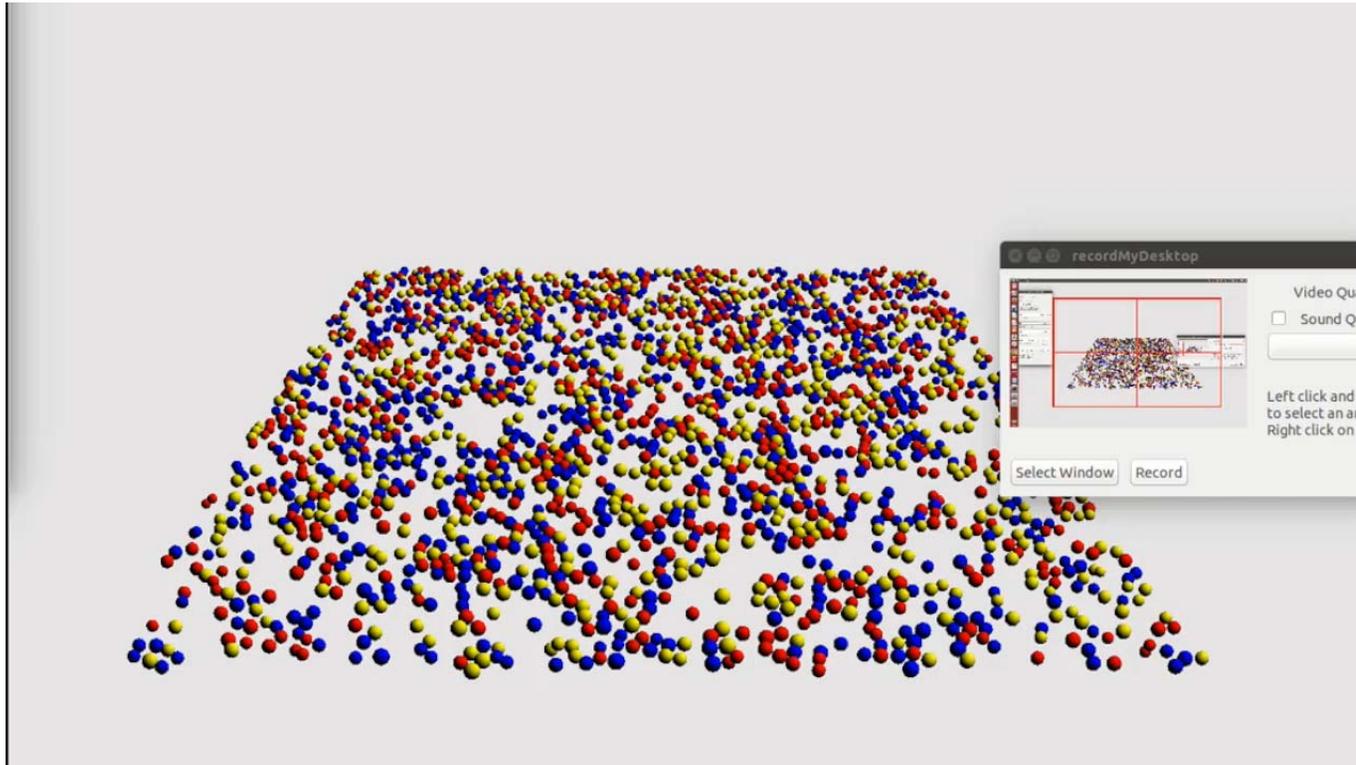


GPUs allow us to maintain high accuracy during the physical interactions within reasonable time constraints (hours vs days)

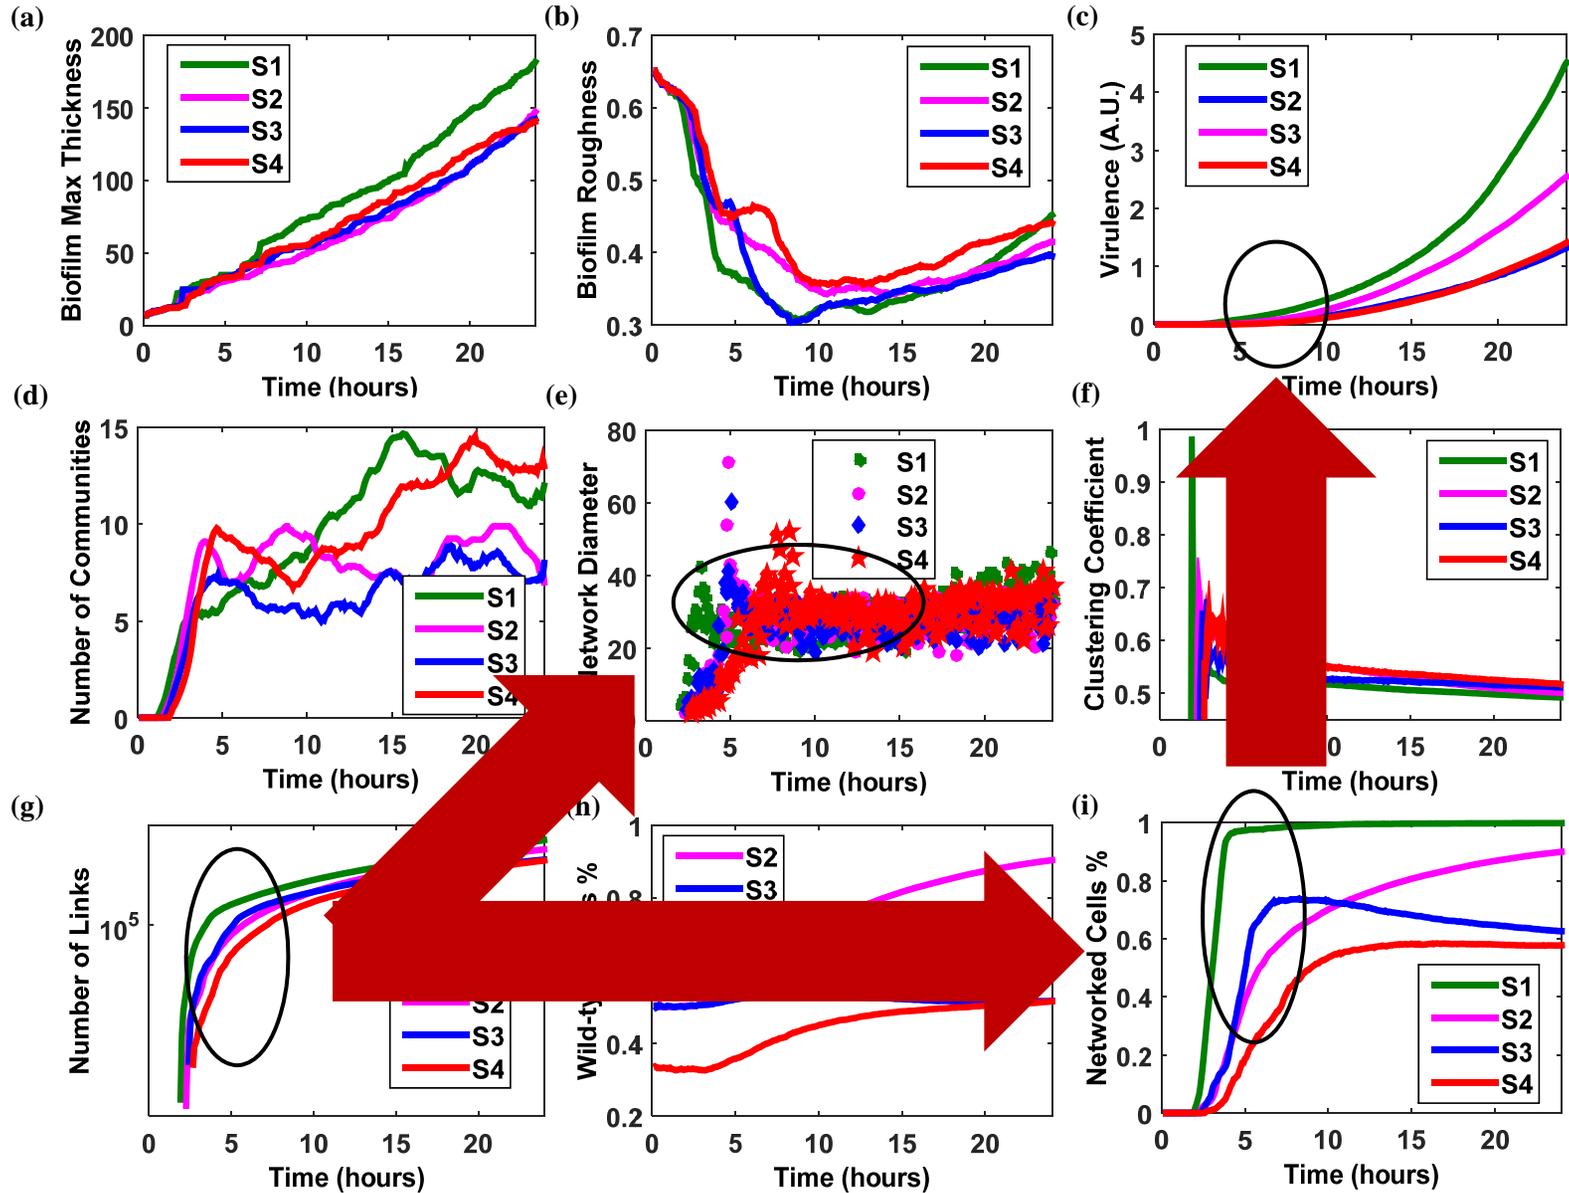
S4 – 1/3WT, 1/3SB, 1/3SN: Bacteria communication enables social intelligence (“molecular tweeting”)

(300x300x500 μm^3)

Red: WT
Blue: SB
Yellow: SN
Green: EPS

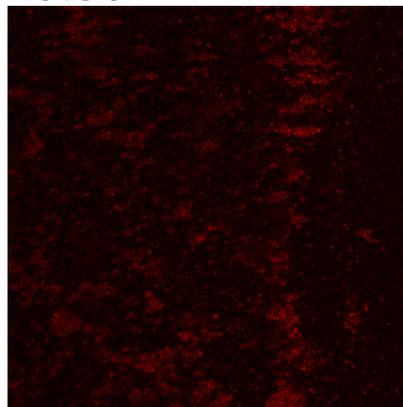
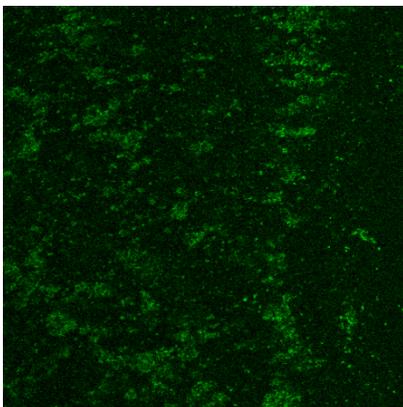


Dynamics of network evolution

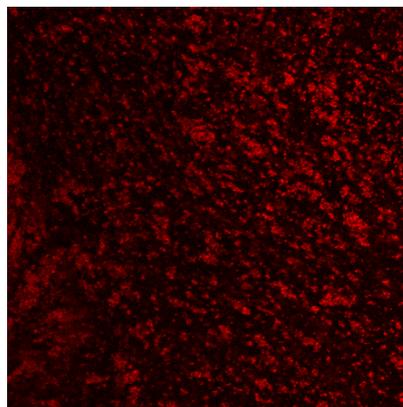
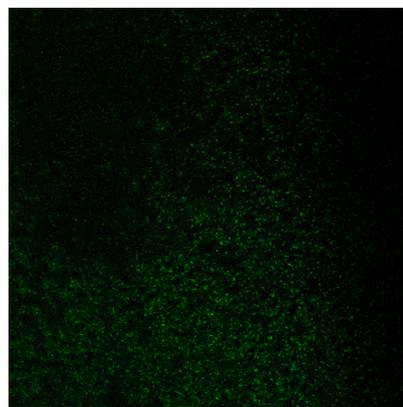


Effect of QSIs

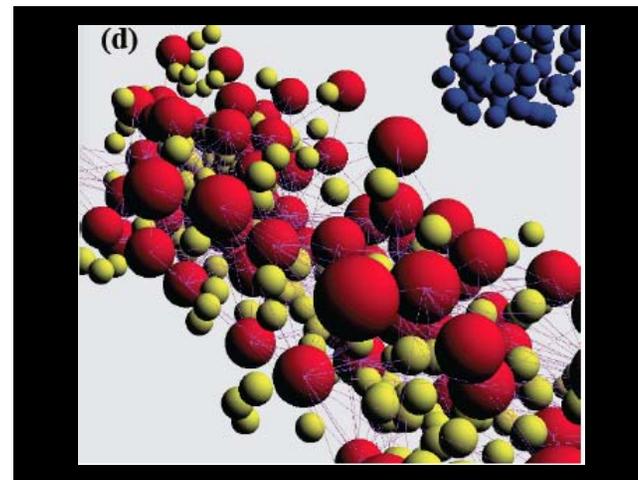
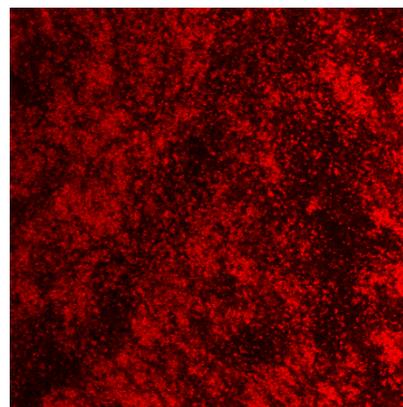
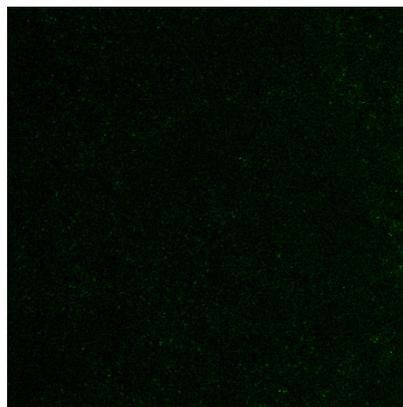
10:30



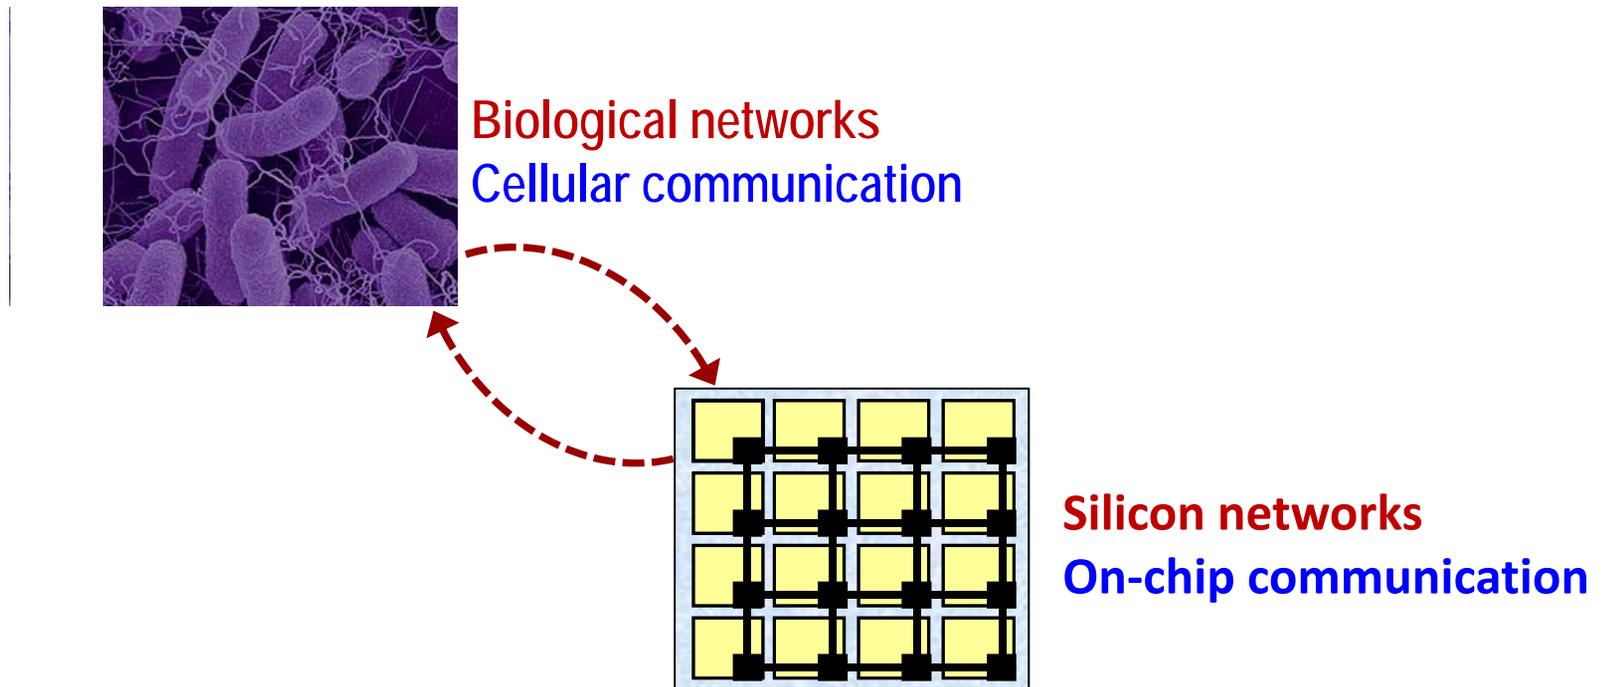
12:47



14:00

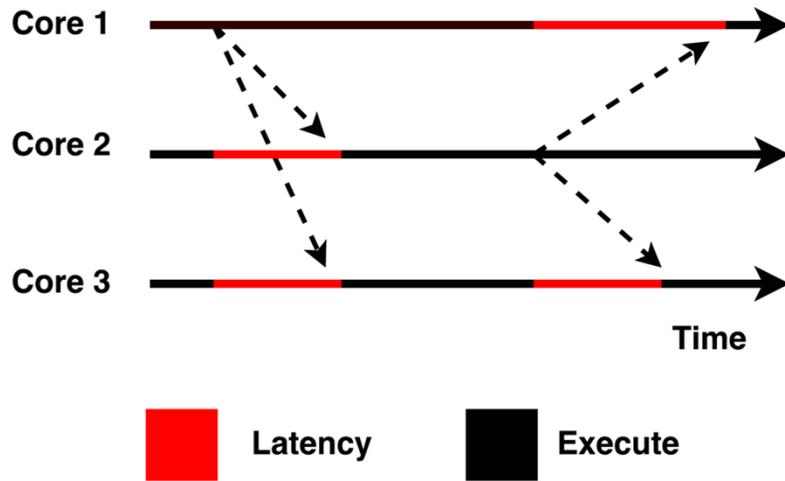


This presentation focuses on the “network” paradigm and ways to engineer communication

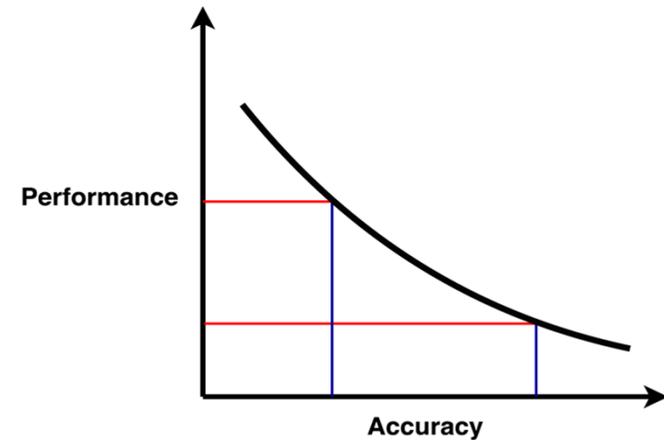


Bacteria form a “social network” to coordinate their behavior, build
22 the biofilm and resist treatment.

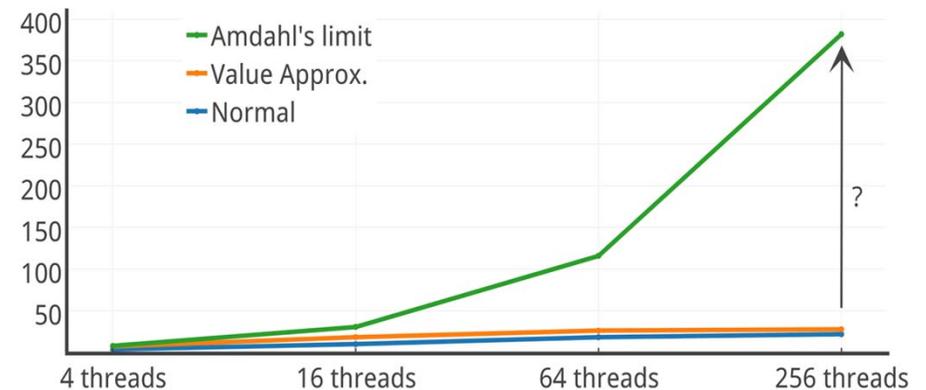
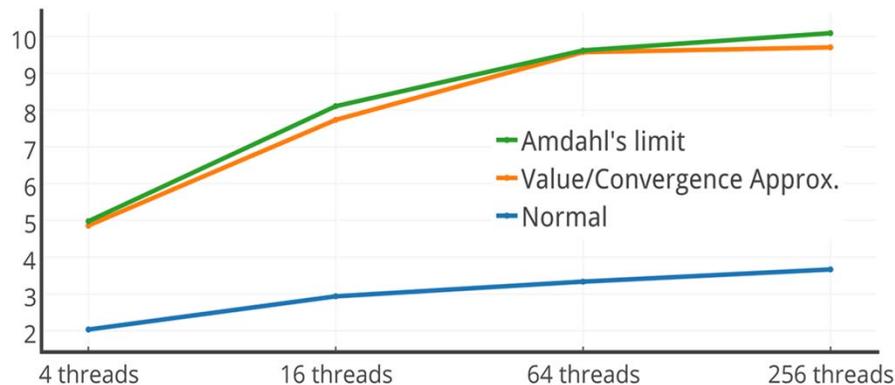
Parallel programs enforce correctness through communication (i.e., synch primitives and cache coherence)



K-Means



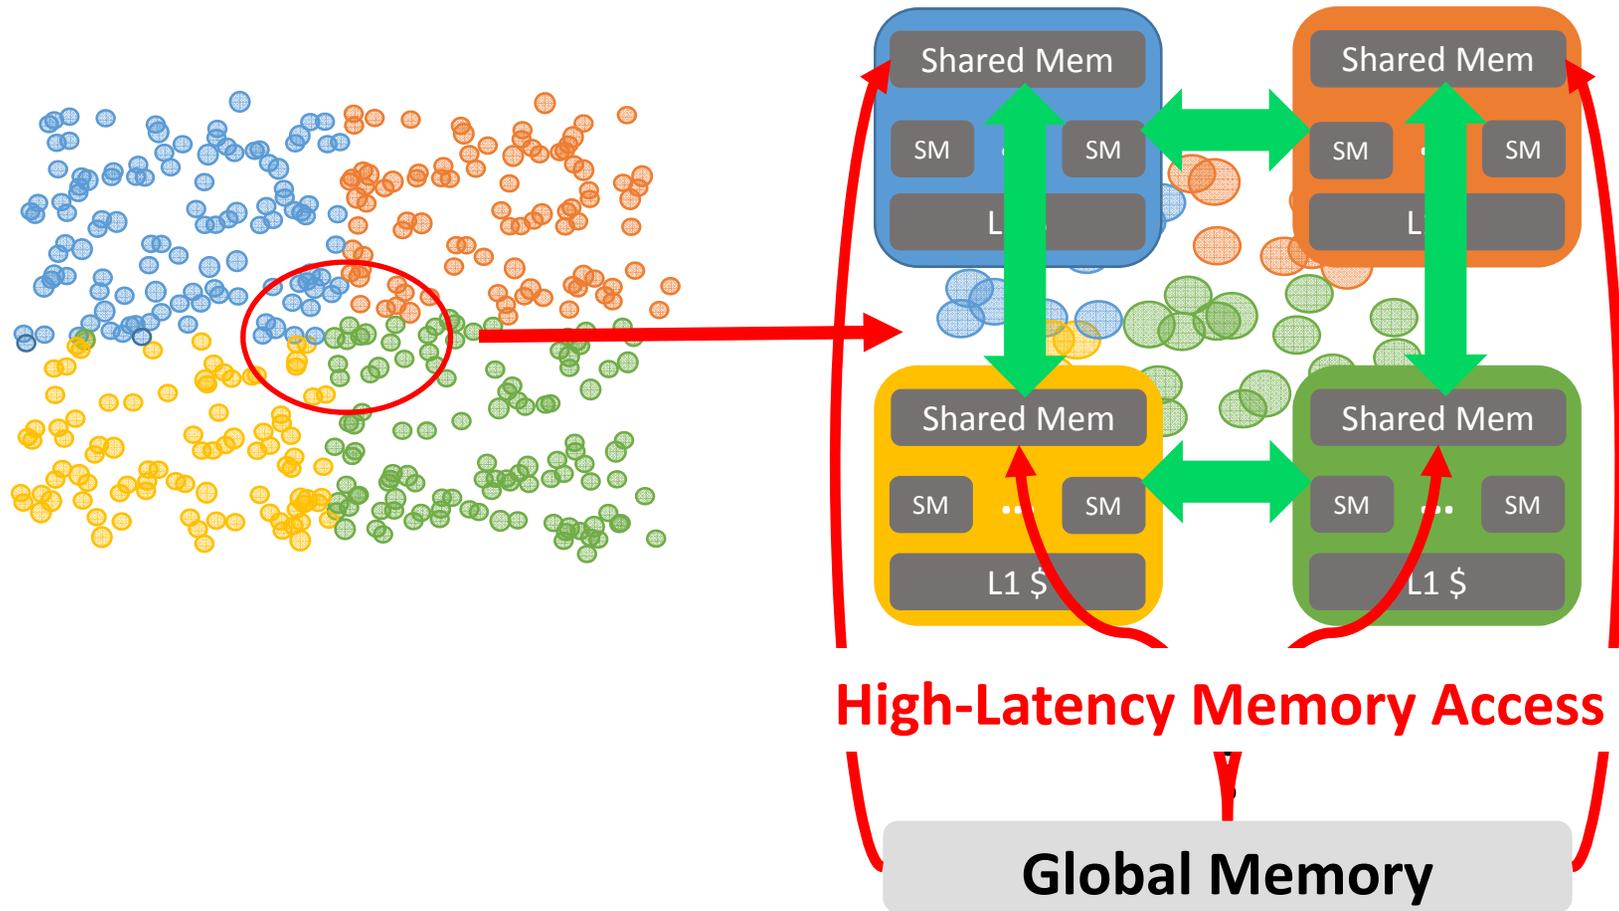
Breadth First Search



[V. Balaji et al., WAX 2016]

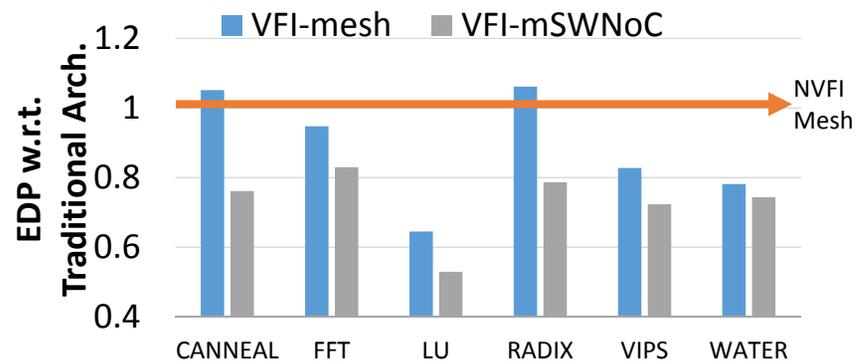
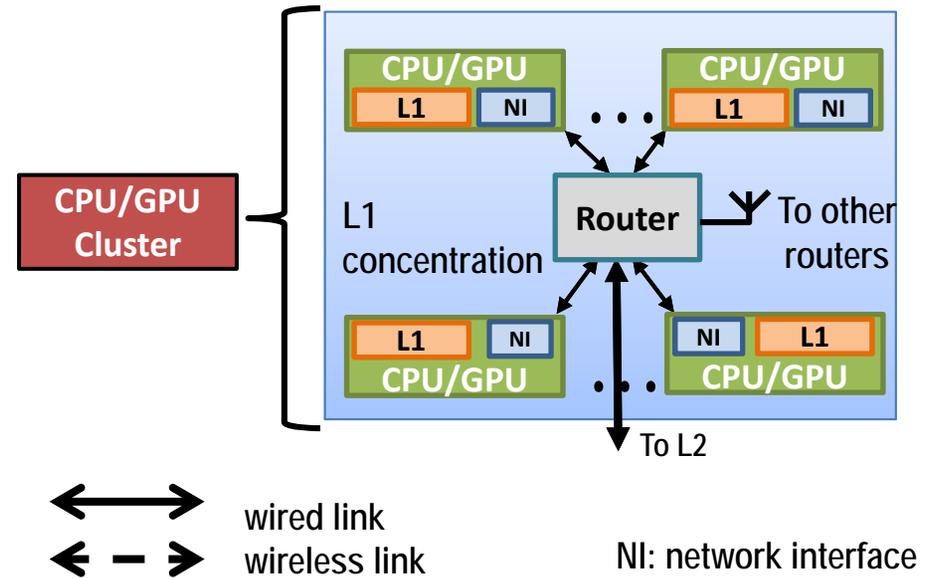
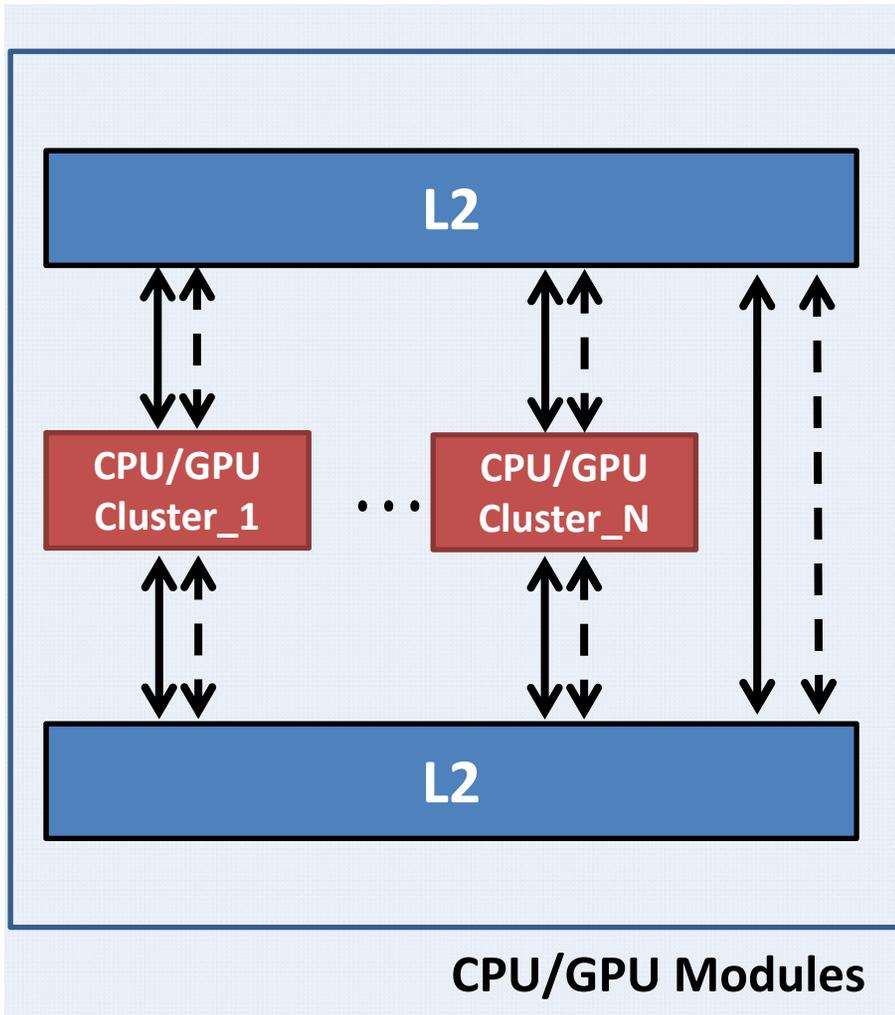
23 **Physical interactions are computationally-intensive and require new approaches to heterogeneous parallel computing.**

Novel distributed CPU-GPU architecture can take advantage of the data-locality between GPU cores



Reduce global access by using low-cost data accesses from adjacent GPUs on an integrated distributed system

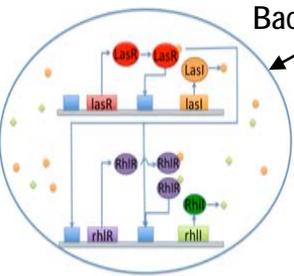
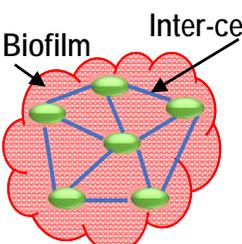
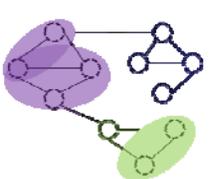
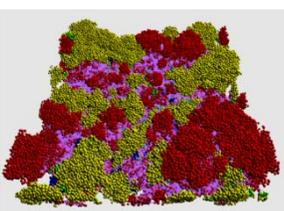
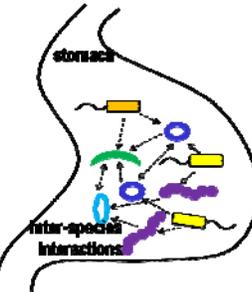
How about future architectures? Can we design an on-chip *Computational Microscope*?

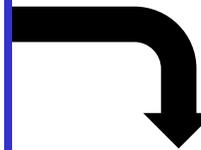


Efficient Wireless NoC Architectures

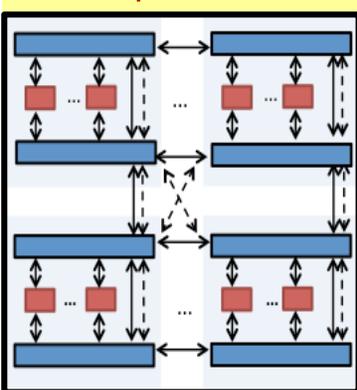
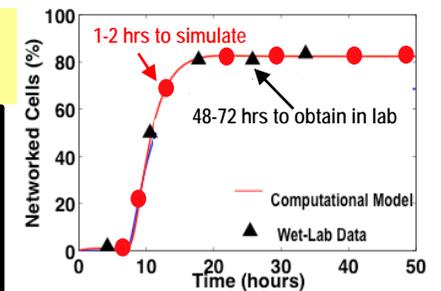
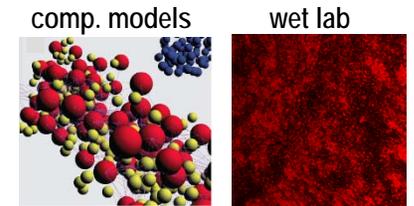
[R. Kim et al., IEEE TC 2015]

There is a lot more to explore...

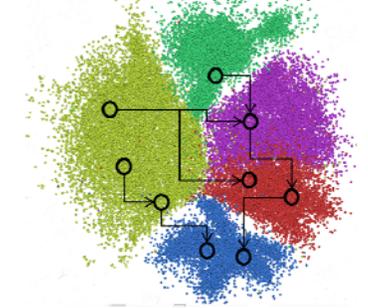
<p>Biofilm Intra-cellular Dynamics</p>  <p>Bacterium (single cell)</p> <p>Need ~ 10 ODEs per bacterium to characterize intra-cell processes</p>	<p>Microbiome Molecular Interactions</p> <p>A: Molecule Protein</p> <p>Similarity</p> <p>Pairwise molec. comparison ~10⁴-10⁶ integer updates per pair</p> <p>A: ...MTVKEQSDIMHGIMSQC... B: ...MTSKEQS--VHDIHSQC...</p>
<p>Biofilm Inter-cellular network</p>  <p>Groups of cells (10³) ~ 10⁴ interaction eqns. (chemical & physical processes)</p>	 <p>Global molec. interaction</p> <p># Sequences: ~10⁶-10¹⁰ # Graph edge traversals: ~10¹⁰-10¹⁵</p> <p>Protein/gene family 1 Protein/gene family 2 Protein/gene family 3</p>
<p>Population of cells (10⁶-10⁹)</p>  <p>Need ~ 10⁷-10¹⁰ eqns for biofilm dynamics (chemical & physical processes)</p>	 <p>Population-scale interaction</p> <p># species per microbiome: ~10²-10⁴ # individual cells: ~10¹⁰-10¹³</p>



SoC-based Computational Microscope

in silico prediction of gene/protein families & molecular pathways



Summary

- **GPU-based simulation is crucial for obtaining faster than real-time results. Only about 1-2 hrs are needed to simulate a scenario that may take 48-72 hours in the lab**
- **Network metrics correlate well with biofilm metrics and can explain population level behaviors of biological significance. Solid basis for studying socio-microbiology...**
- **Inter-cellular communication is crucial for understanding biological systems (which are or should become of interest to computer engineers...)**

Contributors (in no particular order...)

G. Wei (CMU), I. Cazan (CMU), L. Chieh (CMU), K. Bhardwaj (CMU), C. Walsh (CMU), W. Ehrett (CMU), G. Carvajal (CMU), B. Lucia (CMU), L. Hiller (CMU), M. Sitti (CMU), P. Pande (WSU), A. Kalyanaraman (WSU), R. Kim (WSU).

