High Performance Computing Systems (CMSC714)

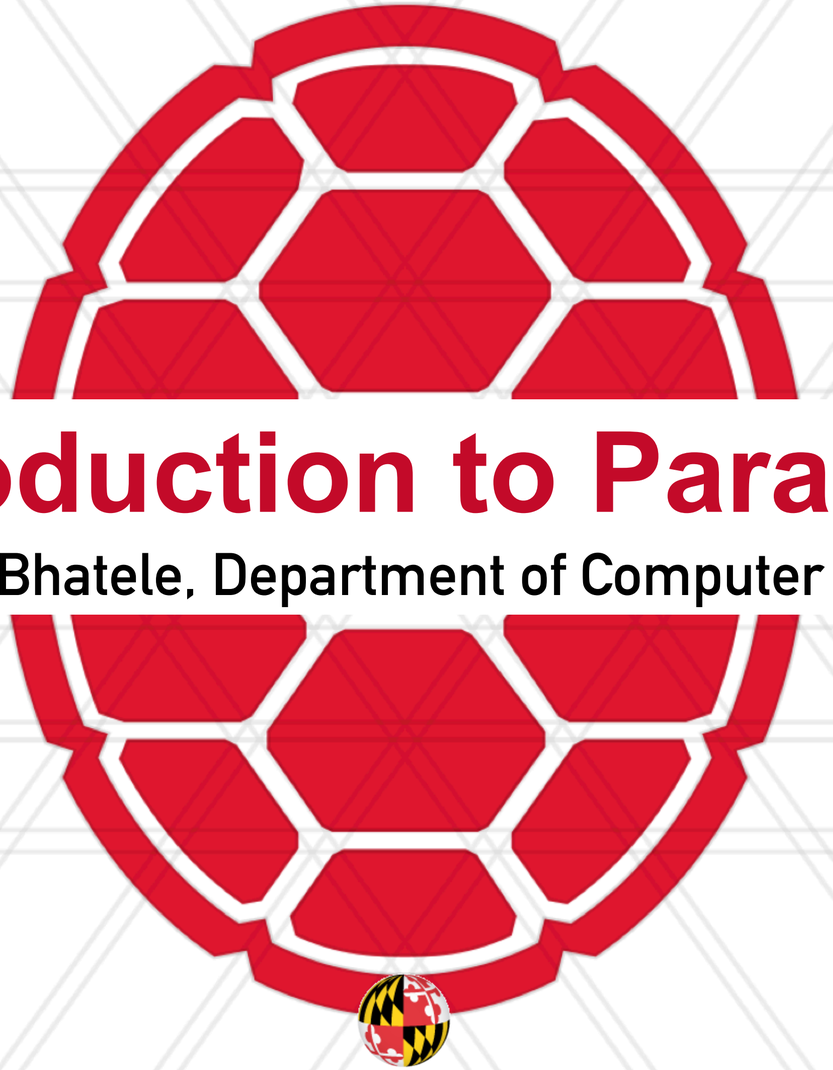

# Lecture 1: Introduction to Parallel Computing

Abhinav Bhatele, Department of Computer Science

UNIVERSITY OF MARYLAND

# A little bit about me …

- Ph.D. from the University of Illinois
- Spent eight years at Lawrence Livermore National Laboratory
- Started at the University of Maryland in August

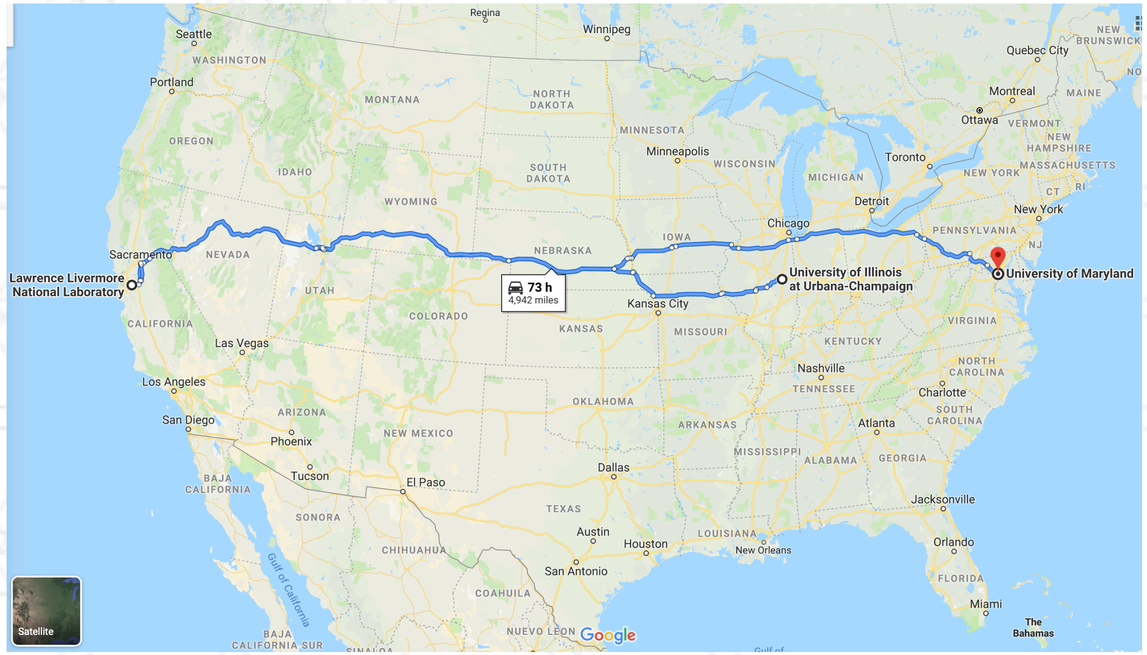

# Introductions

- Name
- MS or PhD / Department
- Area of research
- Why this course?
- Something interesting/ unique about yourself

# This course is

- An introduction to parallel computing
  - Will cover programming models, architectures, tools, systems issues, algorithms and applications
- A qualifying course for MS/PhD
- Work expected:
  - Two programming assignments
  - Mid-term
  - Group project (3 students per group)
  - Classroom participation

# The need for high performance computing

## Drug discovery

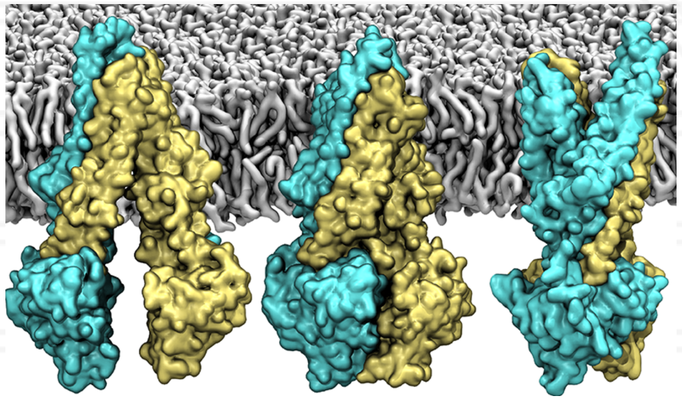

https://www.nature.com/articles/nature21414

# The need for high performance computing

## Drug discovery

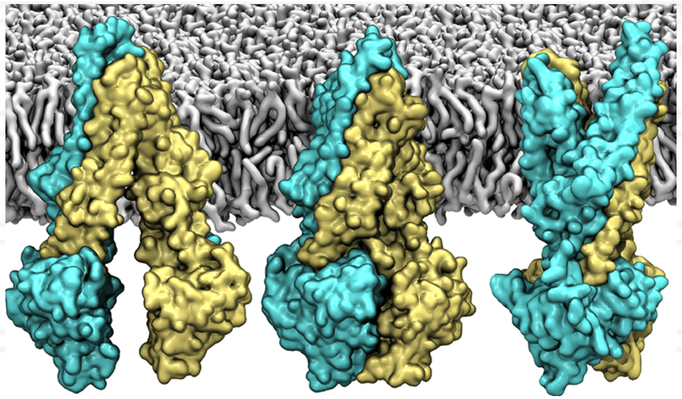

https://www.nature.com/articles/nature21414

## Weather forecasting

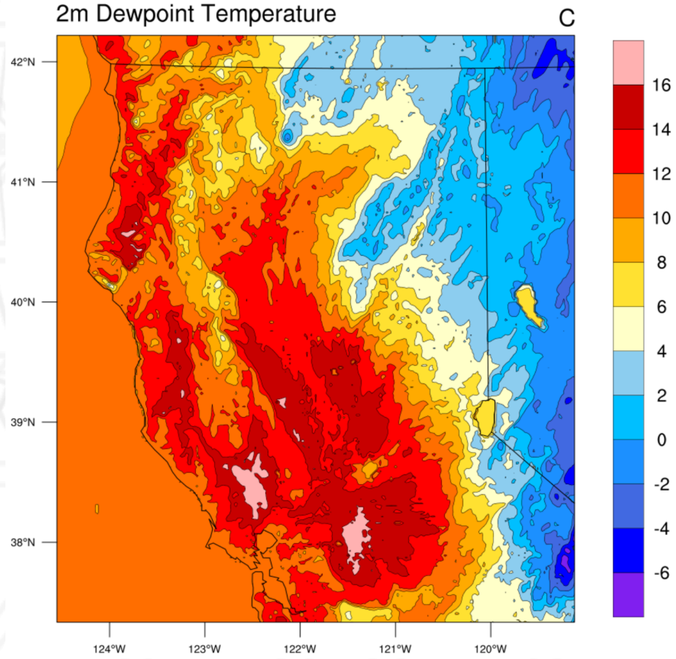


https://www.ncl.ucar.edu/Applications/wrf.shtml

# The need for high performance computing

## Drug discovery

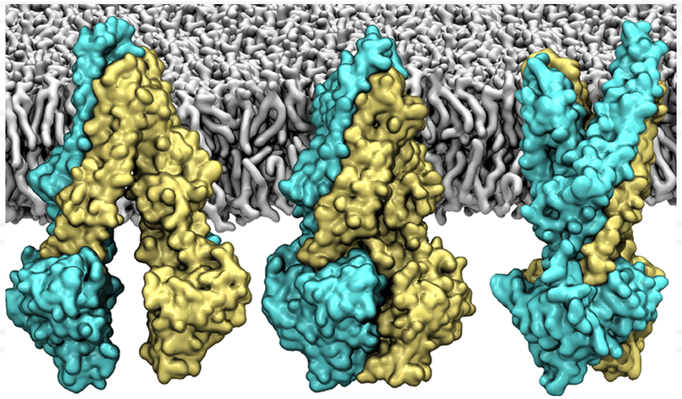

https://www.nature.com/articles/nature21414

## Weather forecasting

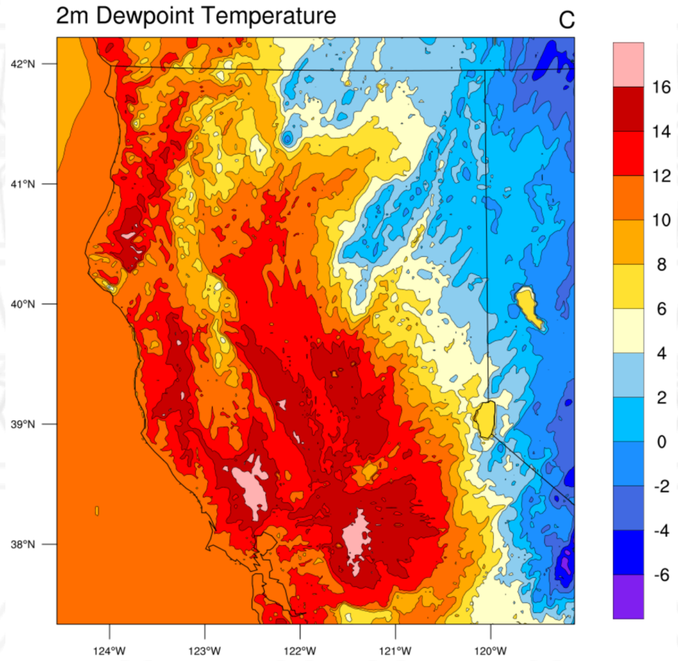


https://www.ncl.ucar.edu/Applications/wrf.shtml

## Study of the universe

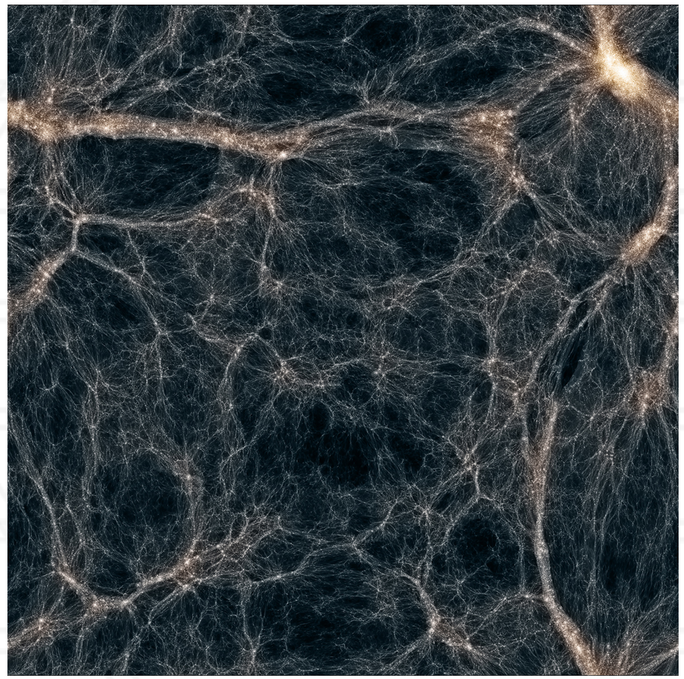

https://www.nas.nasa.gov/SC14/demos/demo27.html

# Why do we need parallelism

- Make some science simulations feasible in the lifetime of humans
  - Either due to speed or memory requirements
- Provide answers in realtime or near realtime

# What is parallel computing?

- Does it include:
  - Grid computing
  - Distributed computing
  - Cloud computing
- Does it include:
  - Superscalar processors
  - Vector processors
  - Accelerators (GPUs, FPGAs)

# Parallel Architecture

- A set of nodes or processing elements connected by a network.

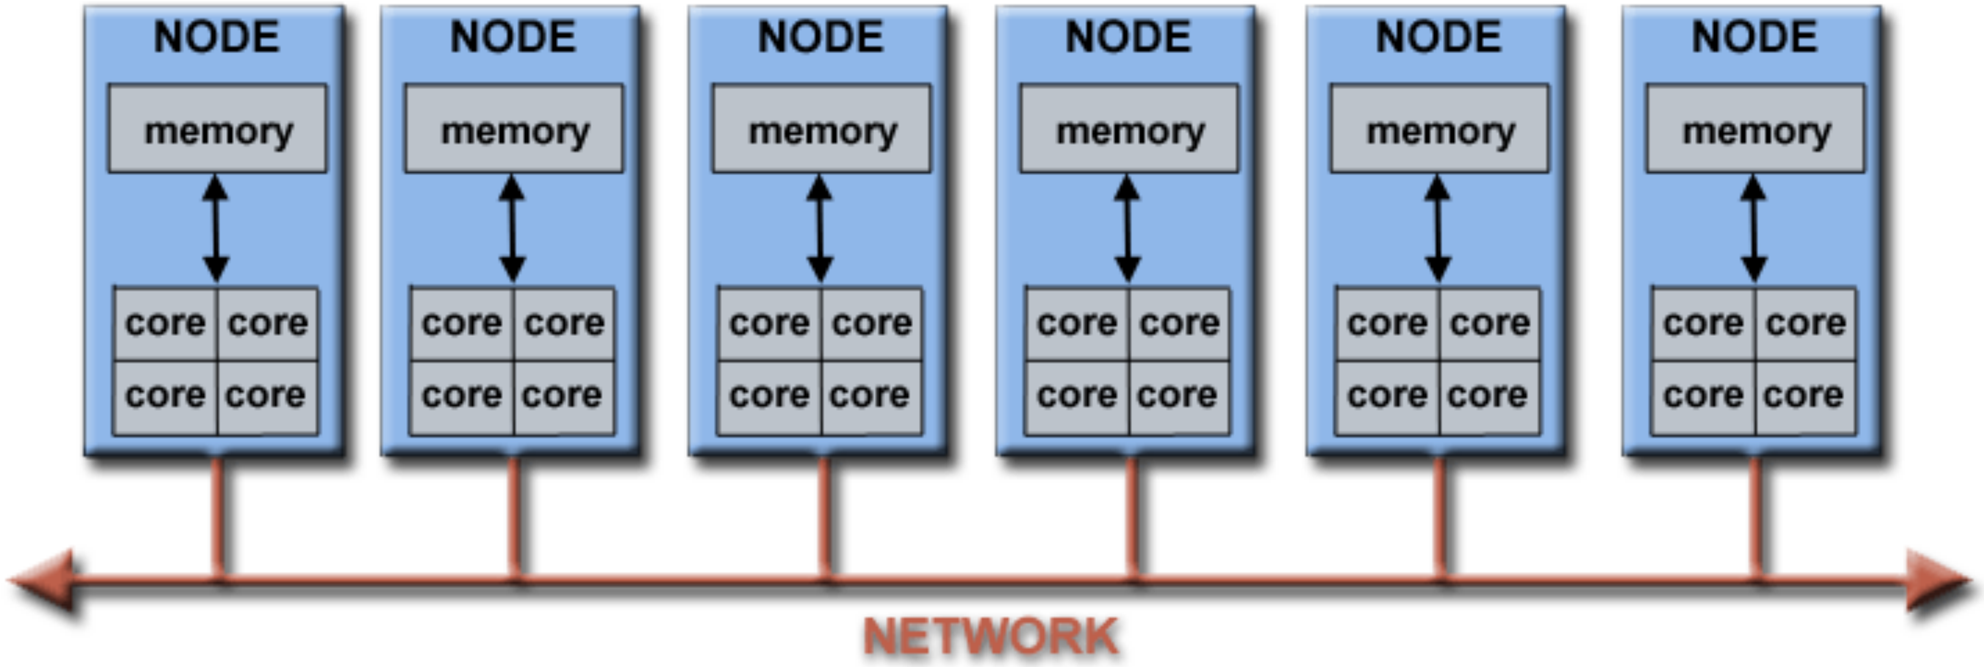


https://computing.llnl.gov/tutorials/parallel_comp

# Interconnection networks

- Different topologies for connecting nodes together
- Used in the past: torus, hypercube
- More popular currently: fat-tree, dragonfly

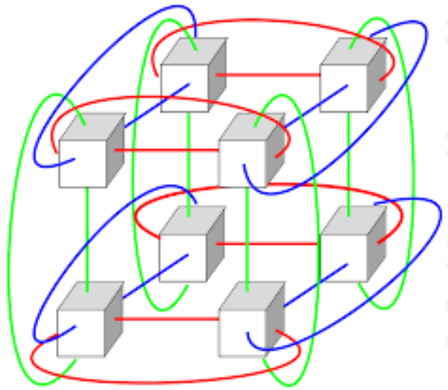

Torus

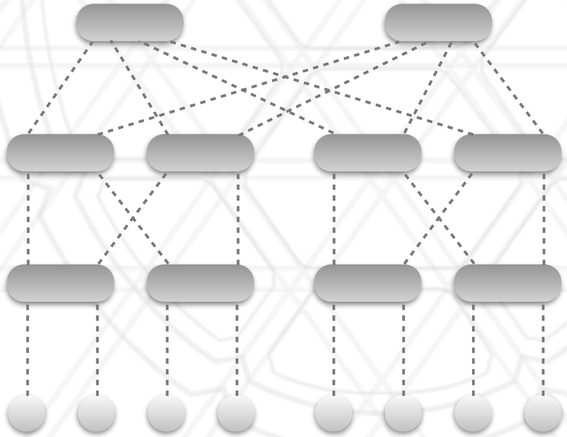

Fat-tree

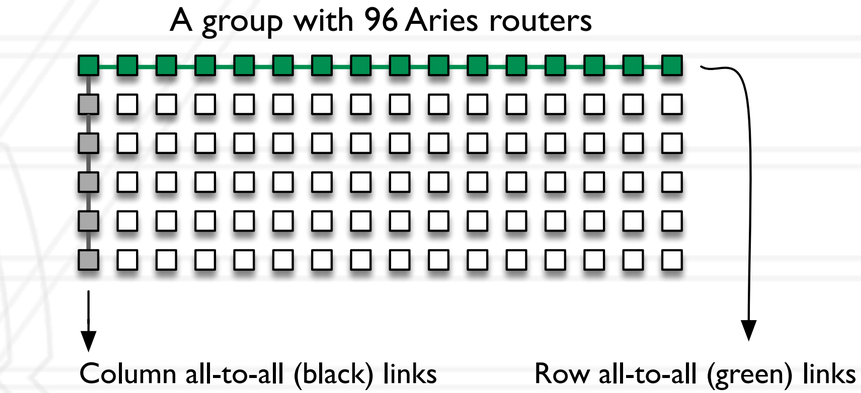


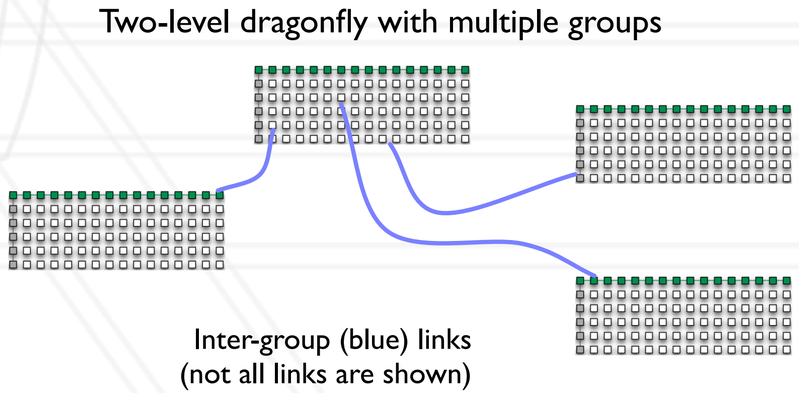


Dragonfly

# Memory and I/O sub-systems

- Similar issues for both memory and disks (storage):
  - Where is it located?
  - View to the programmer vs. reality
- Performance considerations: latency vs. throughput

# System software: Programming models

- Shared memory/ address-space
  - Explicit: Pthreads
  - Implicit: OpenMP
- Distributed memory
  - Explicit: MPI
  - Implicit: Task-based models (Charm++)

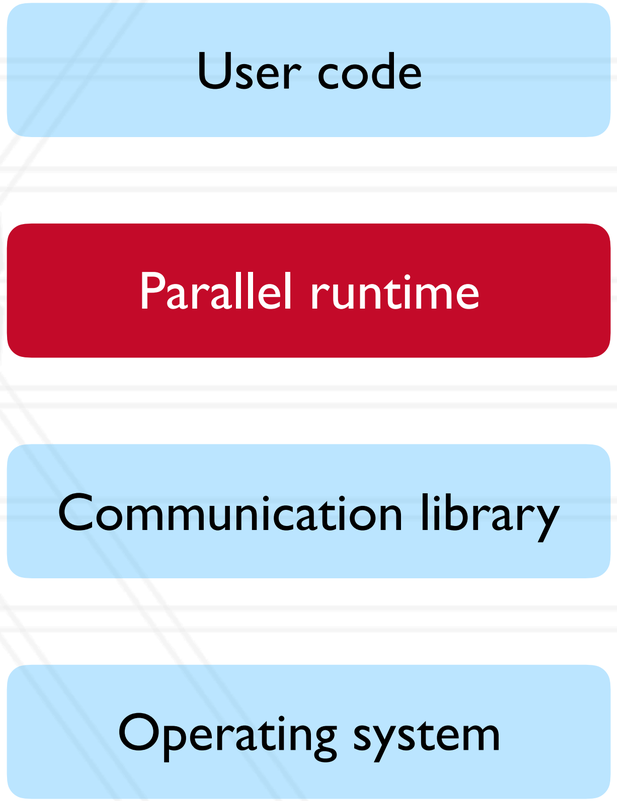

# Performance and debugging tools

- Debugging parallel programs is challenging
- Performance analysis and tuning is critical but hard

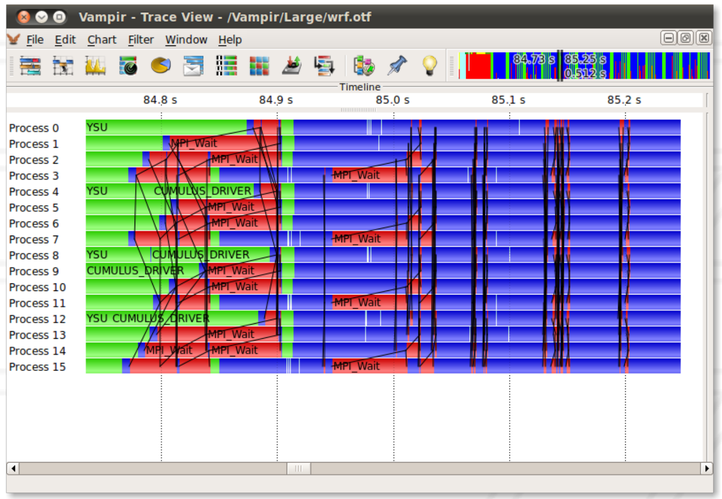


https://vampir.eu/tutorial/manual/performance_data_visualization

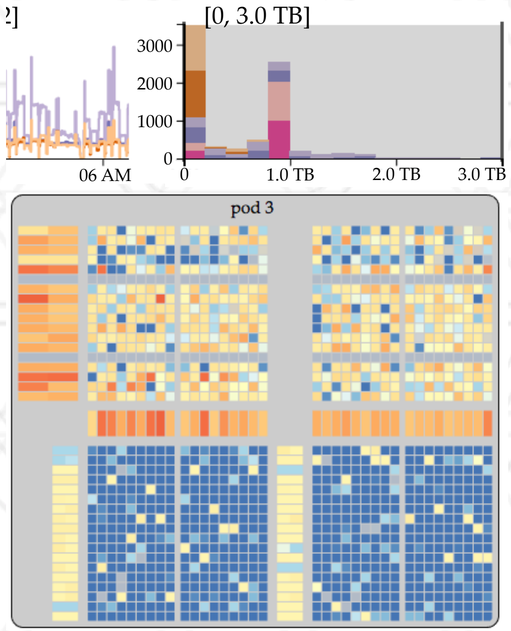


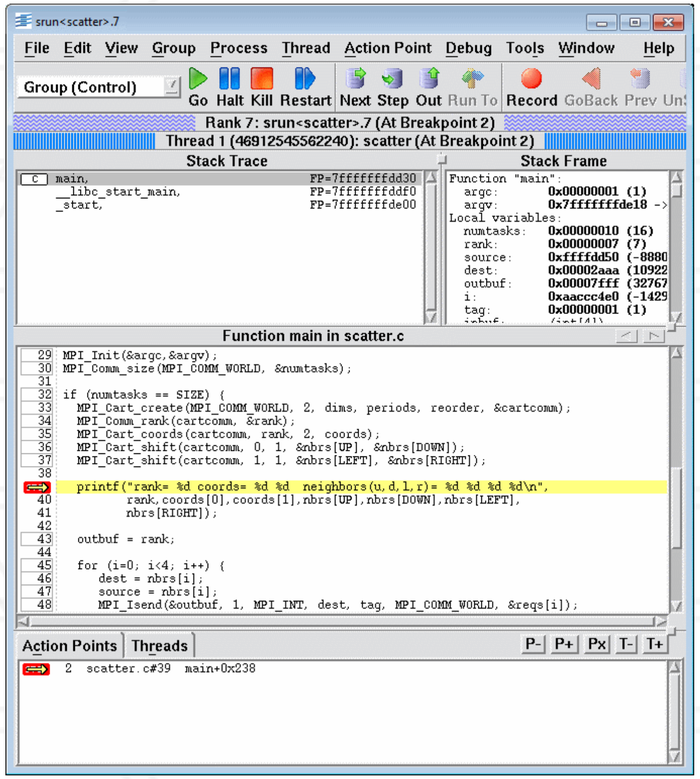


https://computing.llnl.gov/tutorials/totalview/

# Systems Issues

- Operating system noise
- Network congestion
  - Congestion-avoiding routing
  - Parallel I/O
- Job scheduling:
  - Interference from other jobs

# Parallel algorithms and applications

- Parallel Matrix Multiplication
- Parallel Sorting
- N-body calculations
- Discrete event simulations

# Parallel algorithms and applications

- Parallel Matrix Multiplication
- Parallel Sorting
- N-body calculations
- Discrete event simulations
- Molecular dynamics
- Computational cosmology
- Weather and climate modeling
- Discrete-event simulation

# Questions?

UNIVERSITY OF
MARYLAND

Abhinav Bhatele

5218 Brendan Iribe Center (IRB) / College Park, MD 20742

phone: 301.405.4507 / e-mail: bhatele@cs.umd.edu