Introduction to Parallel Computing (CMSC498X / CMSC818X)

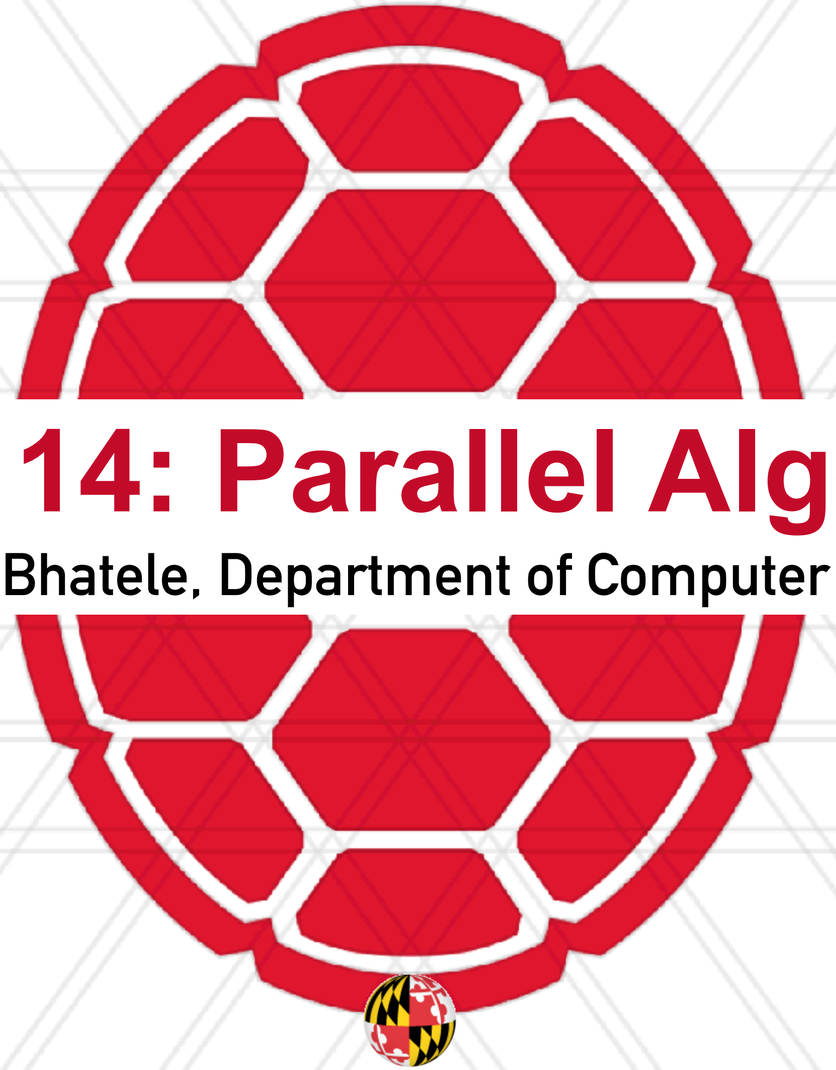

Lecture 14: Parallel Algorithms

Abhinav Bhatele, Department of Computer Science

UNIVERSITY OF MARYLAND

Communication algorithms

- Reduction
- All-to-all

Types of reduction

- Scalar reduction: every process contributes one number
  - Perform some commutative associate operation
- Vector reduction: every process contributes an array of numbers

Parallelizing reduction

Parallelizing reduction

- Naive algorithm: every process sends to the root

Parallelizing reduction

- Naive algorithm: every process sends to the root
- Spanning tree: organize processes in a k-ary tree

Parallelizing reduction

- Naive algorithm: every process sends to the root
- Spanning tree: organize processes in a k-ary tree
- Start at leaves and send to parents
- Intermediate nodes wait to receive data from all their children

Parallelizing reduction

- Naive algorithm: every process sends to the root
- Spanning tree: organize processes in a k-ary tree
- Start at leaves and send to parents
- Intermediate nodes wait to receive data from all their children
- Number of phases: $\log_k p$

All-to-all

- Each process sends a distinct message to every other process
- Naive algorithm: every process sends the data pair-wise to all other processes

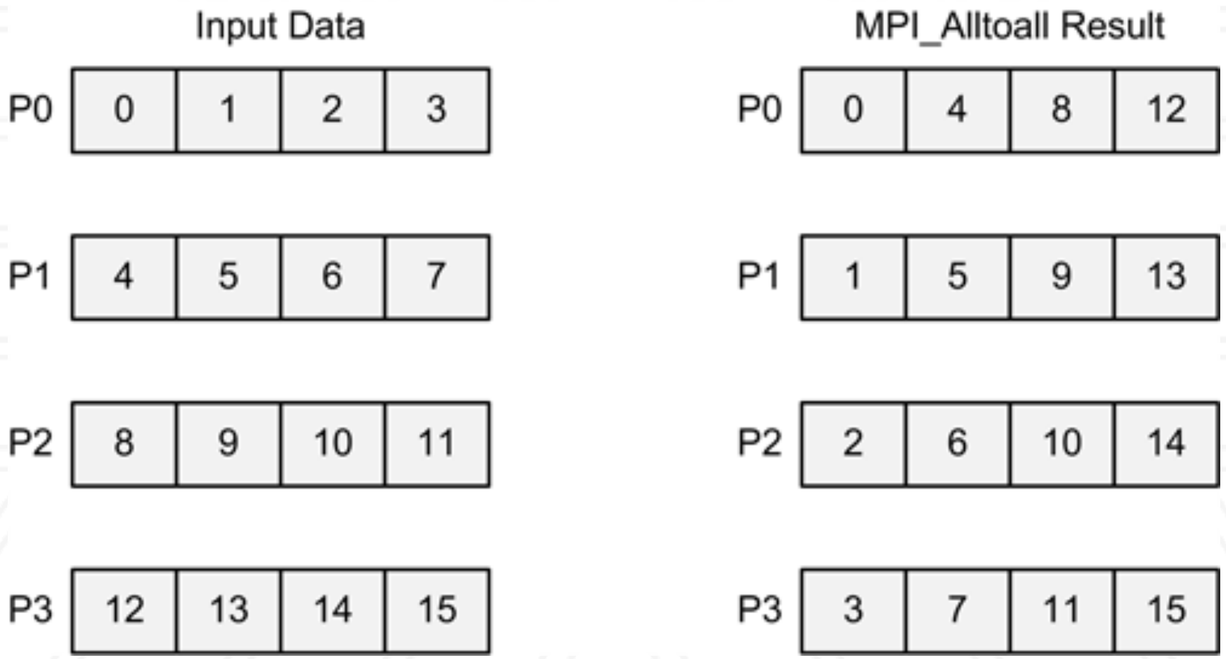

https://www.codeproject.com/Articles/896437/A-Gentle-Introduction-to-the-Message-Passing-Inter

Virtual topology: 2D mesh

- Phase 1: every process sends to its row neighbors
- Phase 2: every process sends to column neighbors

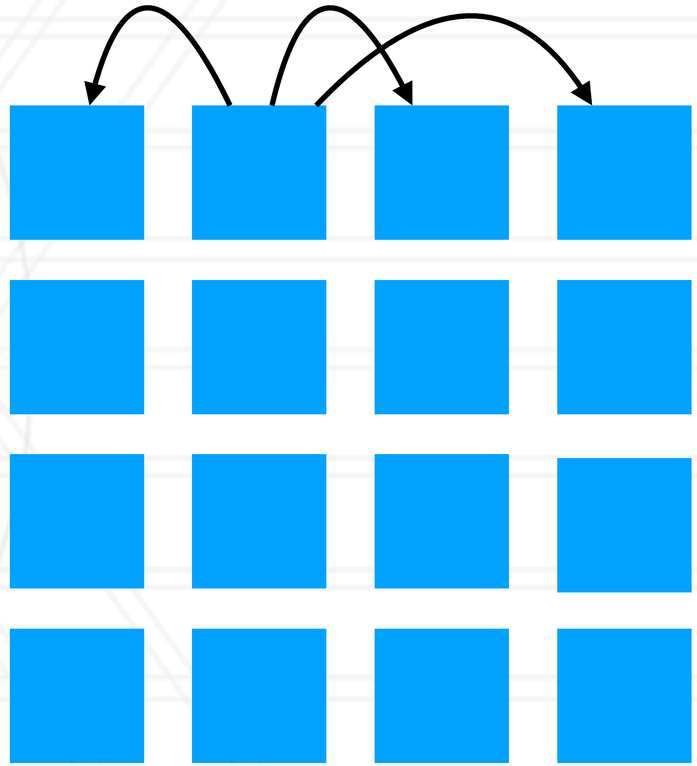

Virtual topology: hypercube

- Hypercube is an n-dimensional analog of a square (n=2) and cube (n=3)
- Special case of k-ary d-dimensional mesh

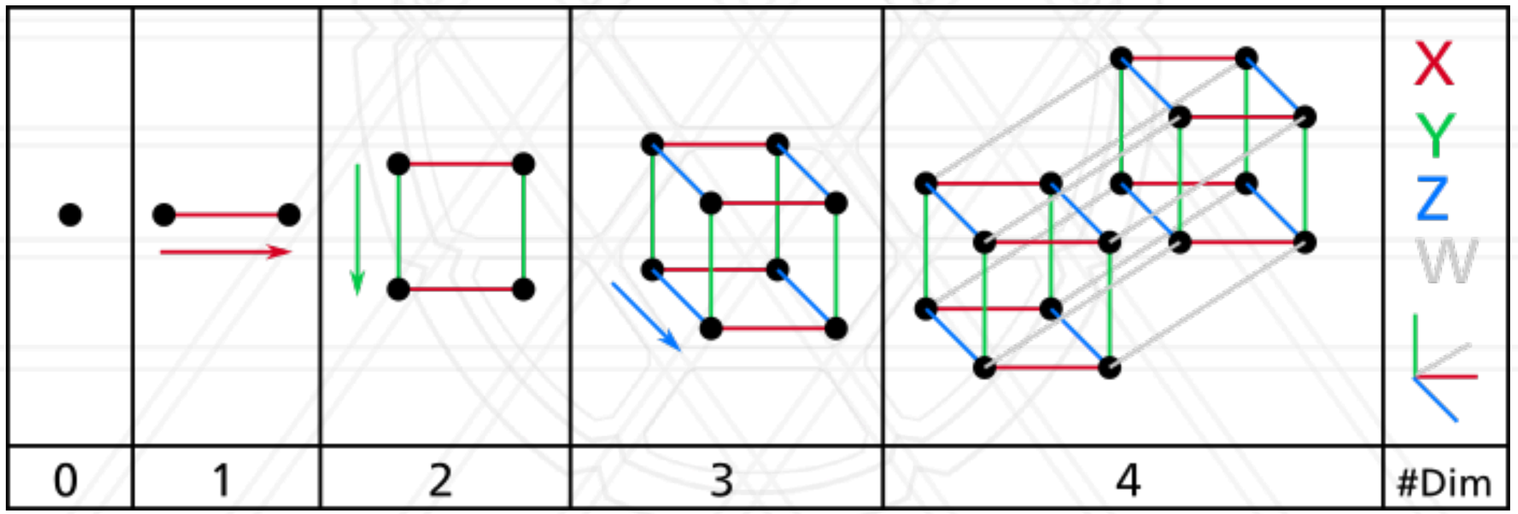

https://en.wikipedia.org/wiki/Hypercube

UNIVERSITY OF
MARYLAND

Abhinav Bhatele

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

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