

COSC 462

Midterm Review

Piotr Luszczek

Parallel Computing and Performance Models

- Speedup: Amdahl's and Gustafson's Laws
 - Be familiar with formulas
 - Look at questions on Piazza
- What's the difference?
 - Parallel vs. concurrent
- Dependencies vs. hazards
- Flynn taxonomy
- PRAM, BSP

Threads

- Thread vs process
 - What is shared what is separate for each?
- Mutex
 - Mutex types
 - Mutex policies
 - How is their performance affected
 - Contested vs. uncontested mutex
- Conditional Variables
- Semaphores
- Atomic instructions

MPI

- Starting and stopping
- Sending and receiving
 - Blocking, non-blocking
 - Standard, Buffered, Synchronous
 - Completion testing and waiting
- Data types
 - Standard-defined
 - User-defined
 - Portable vs. non-portable
- Communicators and Groups
 - Intra vs. Inter
- Collectives
 - Classes

Parallel Algorithm Design

- 4 major phases
 - What are they for?
 - How do they differ?
- Analysis
 - Computation-communication ratio
 - Amount of data coming in/out of a process
 - Total number local process steps
- Know how to apply each phase to basic algorithms from lectures and the textbook
- Consider dimensions in data and process arrangement
 - Apply 1D partitioning to 3D problem
- Consider techniques for dealing with non-divisible dimensions, prime numbers, etc.
 - You are unlikely to resolve all the implementation details on the exam

Problems from Homeworks

- Similar problems may occur on the exam
 - They will be simplified, or
 - The answers will be allowed to be simplified
 - A correct solution may be a combination of faulty code and the discussion of all the problems it has, such as:
 - Scalability problems
 - Lack of parallelism
 - Too much communication
 - Problems with special problem sizes and process counts
- Please provide comments for your code!
 - The code is either correct or has a bug
 - More bug possibilities in parallel
 - Comments communicate to me your intention
 - Almost everybody has good intentions

Textbook Content

- Chapters covered in the exam
 - Chapter 3 Parallel algorithm design
 - Chapter 4 Message passing
 - Chapter 7 Performance analysis
 - Chapter 8 Matrix-vector multiplication
 - Chapter 11 Matrix multiplication
 - Chapter 12 Solving linear systems
 - Chapter 13 Finite difference methods
 - Chapter 14 Sorting
 - Chapter 15 Fast Fourier Transform
- The book covers more than the lectures
 - Stick to topics covered in lectures
 - Use the textbook for additional explanation of lecture topics

General Guidelines

- Show your work
- No calculators
 - Leave large/hard numbers plugged in
 - Estimate the answer if you needed for something else
 - Tell me you used an estimate (but be reasonable!)
 - You can answer based on how you estimated the answer
 - For example when comparing speedup1 with speedup2, you may use the wrong value for speedup1 or speedup2 but the comparison might be valid