Reminders

- Assignment 1 is due April, 15, 11:55PM
  - a well-written, solid explanation of the methodology and rationale pursued to solve the different tasks and the results obtained in your report is essential to pass the assignment
- Mid Semester Exam will be held on Tuesday, March, 26 (tomorrow) during lecture time, remotely on Wattle
  - the weight of the mid semester exam is 10% of your final mark
  - however, it is redeemable with the Final Exam, that is, the combined weight of the MSE and FE will be: max(MSE*10/100+FE*40/100, FE*50/100)

Examination rules and details (I)

click here for the most up-to-date info regarding exam rules

- Exam has to be completed in the course's Wattle page under the "Assessments" section. Please double check that you can access to Wattle before the exam
- Exam is designed to be performed within 4:05PM to 4:55PM. However, -2/+2 minutes tolerance before and after that period will be implemented in Wattle to account for reading time
- You can complete the exam at any venue that has Internet connection (we will not do the exam in the lecture room)

Examination rules and details (II)

- The material assessable by the mid-semester exam includes only the topics discussed until week 5 (inclusive). This spans lectures and labs
- Exam is individual. Finding someone to take the exam on your behalf/help you out during the exam is considered a serious case of academic misconduct
- Exam is open-book. You can use pen and paper
- You can use the Internet to answer the questions. However, AI generative tools such as ChatGPT are NOT allowed. The use of AI generative tools to complete the exam is considered a serious case of academic misconduct

Mid-exam structure

- exam will have several multiple choice questions. Some of the questions might have several correct answers. Wrong answers penalize. Penalization is calculated for each question such that the expected value for a random answer is zero
- the questions are equally weighted
- the questions are designed to assess **your understanding** of the different topics covered along the first 5 weeks (next slides) including your parallel message-passing problem solving and performance modelling skills
- the questions will be similar *in style and difficulty* to some of the previous mid-semester exams, which are all available on the course Wattle page in the dedicated "Past Mid-Semester Exams Papers" section
- in any case, *expect a higher logical reasoning component*

Mid-exam topics (I)

- Introduction (parallelism concept and rationale; scales of parallelism; application areas; Moore's and Dennard's scaling laws; frequency and ILP race; power and ILP walls; multicore revolution; Top 500; current trends in supercomputing)
- Implicit Hardware Parallelism (peak versus effective processor performance; ILP; pipelining; superscalarity; limitations of memory system performance; memory hierarchy; caching)
- Explicit Hardware Parallelism (Flynn's Taxonomy; SIMD versus MIMD parallel architectures; message-passing versus shared address space; UMA versus NUMA; dynamic and static interconnection networks; evaluation metrics for static interconnection networks)
- Message-passing parallelism with MPI (definition; motivation and history; code compilation and execution; point-to-point communication; semantics of point-to-point communication and buffering; deadlocks and deadlock-avoiding techniques; MPI collectives, datatypes and communicators)
- Performance Measures and Models (parallel speedup and efficiency; parallel overheads; hardware, algorithmic, strong and weak scalability; Amdhal's law; Gustafson's law; speed-up and scaled speed-up; timer resolution and overhead)

Mid-exam topics (II)

- Embarrassingly parallel problems (definition; master/slave paradigm; static versus dynamic mapping; performance trade-offs; parallel performance analysis; examples)
- Routing and communication costs (routing, definition and classification; communication cost model, start-up, per-hop and per-word transfer times; routing techniques: store-and-forward and cut-through; modelling the cost of collectives on static interconnection networks; example: one-all and all-all broadcasts; improving speed of communication operations)
- Synchronous computations (definition and examples; static partitioning of grids and matrices; recursive doubling all-gather; performance modelling; the concept of ghost layer aka halo; deadlock avoiding approaches)
- Parallelization by Partitioning and Divide-and-conquer (definition and examples; performance modelling; binary tree-like collective communication)
- Parallelization by pipelining (definition and examples; performance modelling; pipeline latency and average time per problem instance)

Example question

you are given a message-passing code to be executed on a parallel distributed computer with a particular network topology, routing techniques and message-passing libraries, and you are asked to think about:

- communication costs model, and/or
- whether the code is deadlock-free or not, and how to fix it if it may deadlock, and/or
- whether the code produces the correct result, and/or
- whether an alternative network topology would accelerate communication