

Project #2: Parallel Programming Using MPI

Objectives:

- Reason about and design a parallel algorithm in a message-passing environment;
- Experiment with MPI-based programming;
- Evaluate the performance of an MPI program.

Due date: March 3rd, via Digital Drop in BlackBoard.

For this project, you are invited to design and implement an MPI program that counts how many times each word appears in a large text file. Please use the MPI library in C. You may use or not the features introduced in MPI-2, but make sure your decision is based on performance.

Deliverables:

- A document that explains the design of your parallel program and evaluates the performance of your implementation in terms of running time, scalability, and CPU efficiency. Note that your algorithm design might be different than in the first project due to the different architecture (i.e., distributed memory vs. shared memory). Explain your design decisions.

For the performance evaluation you may use other machines than the ones in the C4 lab (if you have access to multi-core machines, for example). However, your code must be able to run on the C4 lab machines for me to test it.

- A README file that shows me how to compile, execute and understand the output of your code. In addition, a makefile would be highly appreciated.
- Your code.