

**CIS 372**  
**Fall 2006**  
**Individual Programming Assignment 3**

**Due Date: Friday, October 6, 2006**

## 1 Objectives

The objective of this assignment is to gain experience in parallel programming with a few different data distributions, and to observe the resulting effect on performance.

## 2 The Problem: The Game of Life

The game of life is a simple cellular automaton where the world is a 2D grid of cells which have two states: alive or dead. At each iteration, the new state of a cell is determined by the state of its neighbors at the previous iteration. This includes both the nearest neighbors and the diagonal neighbors, i.e., like a 9 point stencil without using the value of the cell itself as input. The rules for the evolution of the system are:

- If a cell has exactly two alive neighbors, it maintains state.
- If a cell has exactly three alive neighbors, it is alive.
- Otherwise, the cell is dead.

Your code will need to:

1. Initialize the gameboard
2. Start loop:
  - (a) Print the gameboard (temporarily for correctness checking)
  - (b) Calculate number of live neighbors
  - (c) If (live neighbors = 3) then live
  - (d) If (live neighbors < 2) or (live neighbors > 3) then die
3. End loop
4. Print final gameboard (again temporarily for correctness checking)

## 3 The Assignment

### 3.1 The Basic Sequential Program:

Details of each part follow:

1. **Initialization:** Take the value of  $n$  (the dimension of the board) as an argument to your MPI program. Use array syntax to initialize a board (an  $n \times n$  INTEGER array with value 1 for a live cell and 0 for a dead one) with the following pattern: let  $n = 8$ , and set the row and column nearest the center ( $n/2$ ) to be alive with all other cells dead. For  $n = 8$ , that means that the live cells will be all of row 4 and all of column 4.
2. **Print board:** The board can be printed to the display as lines of 0's and 1's. or as a series of bitmaps which can be animated using xv, by having the board be printed in pgm format (look on the internet for more information on pgm format if not familiar with it).
3. **Update:** Declare a count array the same size as the board ( $n \times n$ ) to contain the number of live neighbors for each point. Include all nearest neighbors, including diagonal neighbors.

4. **Test:** Compile and run the code on a single processor for  $n=8$  and 10 iterations of evolution to check that it produces a set of correct state configurations. Use `xv` to view the resulting arrays to check that your code is working. After checking correctness, you can comment out the print statements inside the loop and keep the final print of the final board.

### 3.2 Parallel Versions:

**Initial Parallel Version with Block Row-wise Data Distribution:** For these parallel versions, you may assume that the number of elements in a row of the square matrix is evenly divisible by the number of processes. Add MPI commands into your sequential code to create a parallel program that distributes the board and count matrices equally in a (ROW-BLOCK) row-wise block data distribution where each process receives and performs computations on a set of consecutive rows of the matrices. Process 0 should print the final board matrix and compare with the results of the sequential program to ensure that correctness is achieved in the parallel version. Test this version of the program on 2, 4, 8, and 16 processes.

**Implementing Additional Data Distributions:** Copy your parallel version of the program and create separate parallel versions that perform data distribution in the following ways:

- ROW-CYCLIC: row-wise cyclic where each row is distributed as a whole row to a given process, and each process is assigned a row in a round robin fashion
- COL-BLOCK: column-wise block distribution where each process receives a set of consecutive columns of the matrix

You should have 1 sequential version and 3 parallel versions of your program. You may use any MPI commands that you see appropriate to accomplish these data distributions.

**Performance Runs:** Increase the size of your arrays to  $n = 2048$ . Compile and run both the sequential and the parallel programs and check correctness one more time by diff'ing the final files. After you are convinced that your parallel versions all work correctly, by checking the printed final matrices against each other, then comment out the printing part of your programs. If you have not already done so, add timing into the program that includes the data distribution, computation, and gathering phases, to be reported as separate timings and then a total time that includes all three phases.

Run each version with  $n = 2048$  with no printing for 1, 8, and 16 processes. Create a table with all of your total timings.

## 4 What to Hand In:

You should hand in the following:

1. **Cover Page:** Title, author, course number and semester, date.
2. **Project Summary and Hypothesis:** First, state and justify your hypothesis about the data distributions. Which data distribution did you expect to run the best, the worst? similarly? Why? Explain your hypothesis in the context of a large number of processes and a small number of processes, and a large matrix versus a small matrix.
3. **Collected Timing Data.** This section includes the tables of the data collected from the performance runs.
4. **Analysis and Conclusions.** Explain your timing results. Describe your observations concerning the different data distribution methods, and how it compares to your hypothesis. Explain why you believe the timings are different or similar between the methods. How does it compare to your hypothesis? Discuss possible reasons for inconsistencies or discrepancies in your data versus what you would have expected to happen.
5. **Appendix:** Your code for the sequential version and each of the parallel versions. Be sure to label each version clearly as `seq.c`, `rb.c`, `rc.c`, and `cb.c`. A script of a run with  $n=8$ , our initial configuration as described above and printing of the output to the screen, for 5 iterations.

Please staple all parts of your lab together, and label each piece. Be prepared to discuss your results in class.

## 5 Criteria for Evaluation

Your lab will be evaluated according to the following criteria:

1. 13 pts: Sequential game of life
2. 15 pts: Parallel version 1: row-block
3. 15 pts: Parallel version 2: row-cyclic
4. 15 pts: Parallel version 3: column-block
5. 15 pts: Table of all timings
6. Report
  - (a) 2 pts: cover page
  - (b) 10 pts: hypothesis and project summary
  - (c) 15 pts: Analysis and discussion

## 6 Extra Credit - 6 pts

Modify your program to work correctly when the number of elements in a row of the square matrix is not evenly divisible by the number of processes. Show that your program works correctly by running each parallel version with a matrix size that fits this description. Hand in the code and the output of the runs.

## 7 Notes on Multidimensional Arrays

To be able to access consecutive rows of a 2-dimensional array in C, you must allocate space for the entire matrix together with one malloc or statically. Otherwise, it is possible that separate malloc's get memory allocated from nonconsecutive locations.