

Parallel Game of Life

Mingyuan Ding (mingyuad), Li Shi (lishi)

Project Webpage

15618-parallel-game-of-life.github.io

Summary

We plan to implement a parallel Conway's game of life simulation and a reinforcement learning AI model to play against the player or itself.

Background

The Game of Life is a cellular automation devised by the British mathematician John Horton Conway in 1970. The evolution of cells is determined by their initial state on a two-dimensional grid, where each cell can either be alive or dead. Our project can be divided into two phases. The first phase includes implementing a parallel game simulator that supports the traditional rule of the game of life, i.e.,

1. Any live cell with fewer than two live neighbors dies, as if by underpopulation.
2. Any live cell with two or three live neighbors lives on to the next generation.
3. Any live cell with more than three live neighbors dies, as if by overpopulation.
4. Any dead cell with exactly three live neighbors becomes a live cell, as if by reproduction.

In the second phase, we will try to modify the game to have two players and play in a competitive fashion. Two players place initial cells on the grid with given constraints. A possible modification could be that any cell's next state is determined by the delta of the number of friendly live neighbors and enemy live neighbors. We plan to train an AI using reinforcement learning to play against either human players or itself. The training and execution of AI can benefit from parallel programming.

Challenges

In the first phase, the challenge is to search and reach the roof of the machine we are running on. As the game of life is a communication-intensive program, it requires efforts how to properly assign and schedule tasks on different computing units and perform extra optimization to accelerate large-scale game simulation.

In the second phase, selecting and constructing the model can be challenging. Despite the simple rules, the game of life can take a long time to settle into a stable state and is very sensitive to the initial state. As the game involves long-term planning and strategy, the model also needs to perform searching in a large action space and it is challenging to perform optimization with parallelism in this part.

Resources

We are not planning to use any starter code but we are going to do a literature review and pick a reinforcement learning model from recent papers. A reinforcement learning framework may be used to achieve better training results. We are planning to select a parallel programming model like OpenMP or CUDA for simulation and use CUDA for reinforcement model training. GHC machines will be our main platform.

Goals & Deliverables

Plan to achieve (Project 1st phase):

1. A sequential version of the game of life simulation as infrastructure
2. A parallel version of the game of life simulation with qualitative analysis of speedup
3. Visualization of the game of life

Hope to achieve (Project 2nd phase plus optimization):

1. Implementation of reinforcement learning to play the multiplayer version of the game of life
2. Further runtime analysis and optimization of parallel simulation
3. Optimization of training and execution of reinforcement learning model

Platform Choice

We are going to accelerate the simulation of the game with either OpenMP or CUDA, and train and run the AI model using Python or C++ with CUDA. Since we are planning to use both the CPU and the GPU, GHC cluster is the best choice.

Schedule

Week 1 (4/2 - 4/8): Review course materials. Study background of Conway's game of life and do literature review.

Week 2 (4/9 - 4/15): Implement a sequential version of game simulation. Select a proper parallel programming model and start to implement a parallel version of game simulation.

Week 3 (4/16 - 4/22): Visualize the game simulation. Implement a basic reinforcement learning model for the multiplayer version of game of life.

Week 4 (4/23 - 4/29): Fix potential problems in the parallel version of simulation and optimize performance. Improve the performance of the reinforcement learning model.

Week 5 (4/30 - 5/5): Do data analysis and prepare for the final poster presentation.