Computational Details

Grid: 200x200 = 40,000 states

We have one CPU with one GPU

Policy Iteration algorithm

For each state, the computer computes the value of going to each of the other 40,000 states.
This is done for all 40,000 states.

Then, when it has the 40,000 new policies, it solves the 40,000 linear equations that define the value of that policy.

That is the new value function and new policy.

Computer run

Send all states and current value function to GPU
    3500 gpu cores
    For each state, a gpu core computes values of all possible transitions,
        finds max, and
        tells you the policy and max value
    When gpu has all the solutions, it computes the value of the new policy

Time to solve a single problem: ONE MINUTE
(If used a solver, time would be about 20 minutes)

GPU is great for evaluating formulas, such as my closed-form solution.
GPUs will not take a serious solver.

In our problem, CPU does nothing other than sending work to the GPU

Hardware

Google CoLab

Free if you don't use it much

We only use the $10/month

Uses Python and numba

GPU is straightforward to use with this setup

Could have multiple gpu