Massively Parallel Model of Evolutionary Game Dynamics

Amanda Peters Randles

School of Engineering and Applied Sciences, Harvard University.

Summary

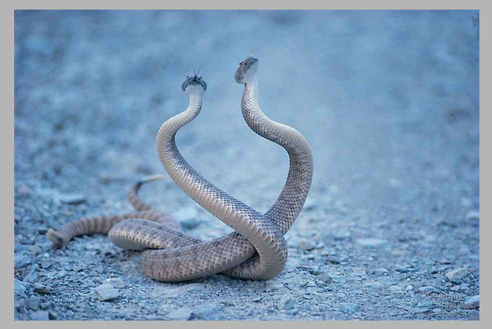

To study the emergence of cooperative behavior, we have developed a scalable parallel framework. An important aspect is the amount of history that each agent can keep. When six memory steps are taken into account, the strategy space spans 2^{4096} potential strategies, requiring large populations of agents. We introduce a multi-level decomposition method that allows us to exploit both multi-node and thread-level parallel scaling while minimizing the communication overhead. We present the following contributions:

- A production run modeling up to six memory steps for populations consisting of up to 10^{18} agents, making this study one of the largest yet undertaken.
- Results exhibiting near perfect weak scaling and 82% strong scaling efficiency up to 262,144 processors of the IBM Blue Gene/P supercomputer and 16,384 processors of the Blue Gene/Q.

Our framework marks an important step in the study of game dynamics with potential applications in fields ranging from biology to economics and sociology.

Population Model

We focus on 2-player Prisoner's Dilemma.The two agents can choose to either cooperate (C) or defect (D). The reward table is below:

	Player 1	Player2
Both players cooperate	3	3
Player 1-C Player 2-D	0	4
Both players defect	1	1
Player 1-D Player 2-C	4	0

Cooperative behaviors may be exhibited because in real-life situations opponents have an expectation of playing an opponent again in the future. In such a repeated game, the two agents face each other in many different *rounds*. Each agent's fitness is assessed by summing the fitness achieved in each round of game-play. The goal of each agent is to accumulate the highest fitness possible. Each agent can use historical information to determine their best move and to maximize their fitness. The rules that determine an agent's next move are referred to as a *strategy*.

Memory Steps	Number of Strategies
1	2^{4}
2	2^{16}
3	2^{64}
4	2^{1024}
5	2^{2048}
6	2^{4096}

Number of Pure Strategies per memory step.
Red Box: Previous simulation capability

A *pure strategy* gives a complete definition of how the agent will proceed in each round. Given state x, agent will always take a set action.

SSets

Our approach considers three major entities:

- Agents
- Strategy Sets (Ssets) -- agents assigned the same strategy
- Nature Agent

This enables a tiered parallel scheme whose hierarchy is depicted in the figure below.

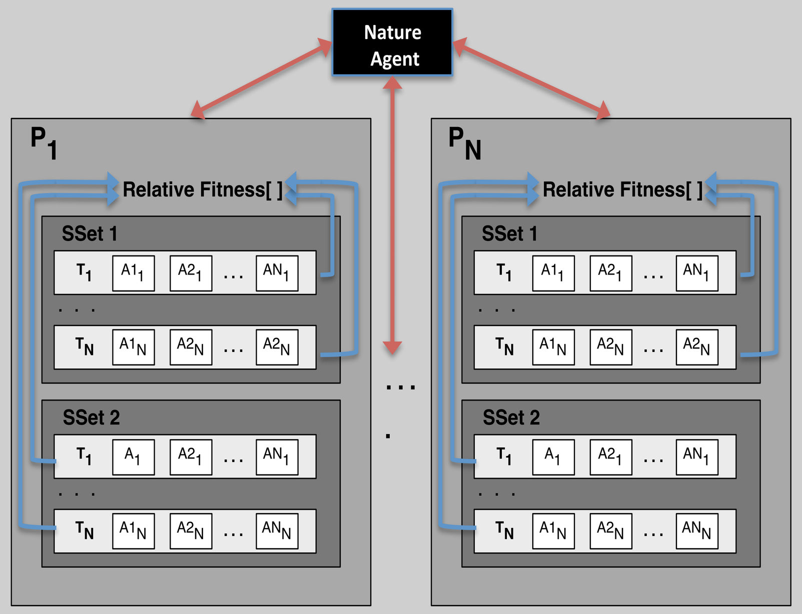

- Each grey box shows the Ssets on each processor.
- Within each Sset, there is a set of agents split across separate threads.
- All threads update a local array of the relative fitnesses of Ssets on that processor.

Non-trivial Communication Pattern:

- At random generations, the agents are all interrupted by the Nature Agent's random global broadcast of SSets for pairwise comparison (PC) or mutation.
- Point-to-point communication between processors in an Sset.
- Broadcast of new overall population state.

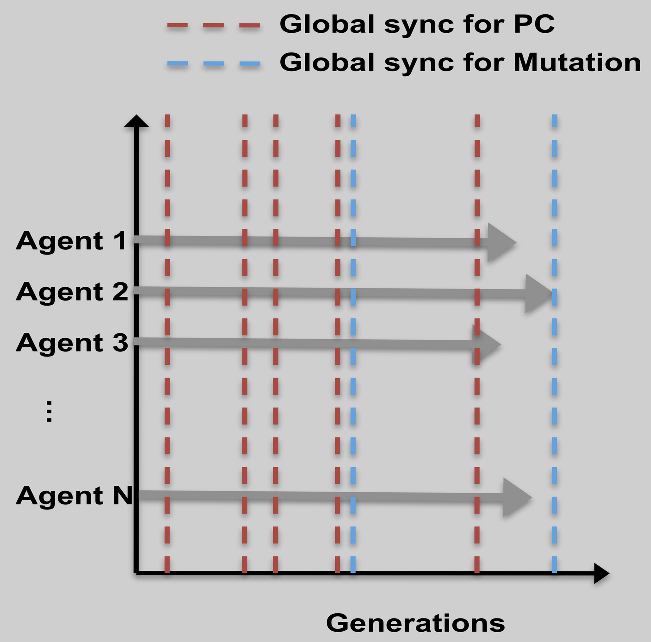

Small-scale Studies

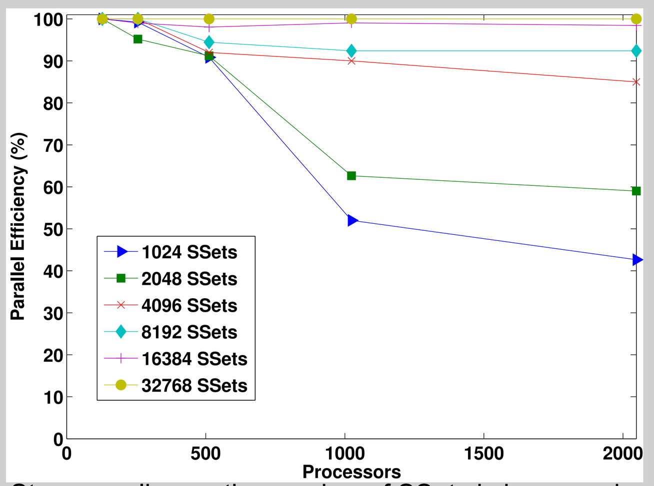

Strong scaling as the number of SSets is increased.

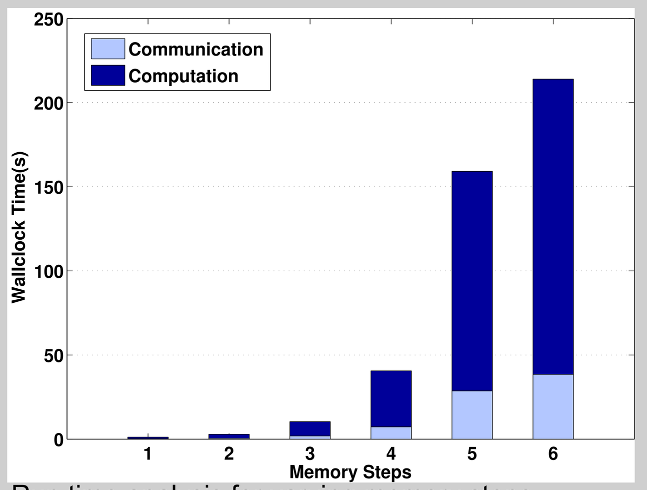

Run time analysis for varying memory steps.

- With fewer than 4,096 Ssets per proc, communication overhead decreases the parallel efficiency.
- Taking into account a greater number of memory steps increase the runtime but has little impact on parallel efficiency

Large-scale Studies

Weak Scaling: 4,096 Ssets per processor
Maximum Population Size: 1,073,741,824 Ssets
Memory Step Level: **6**

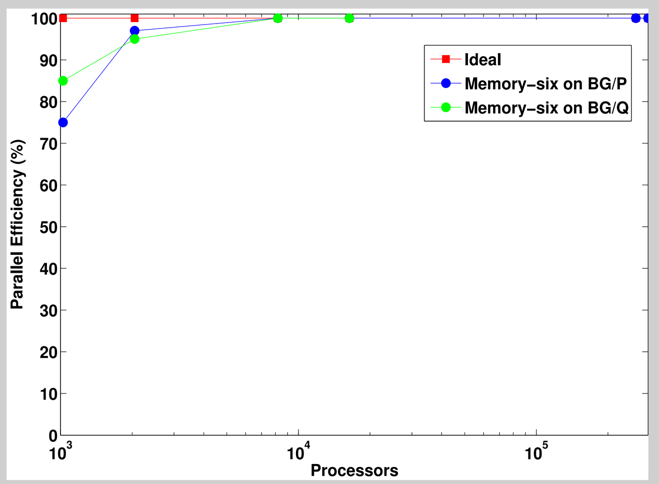

Weak Scaling up to 294,912 processors.

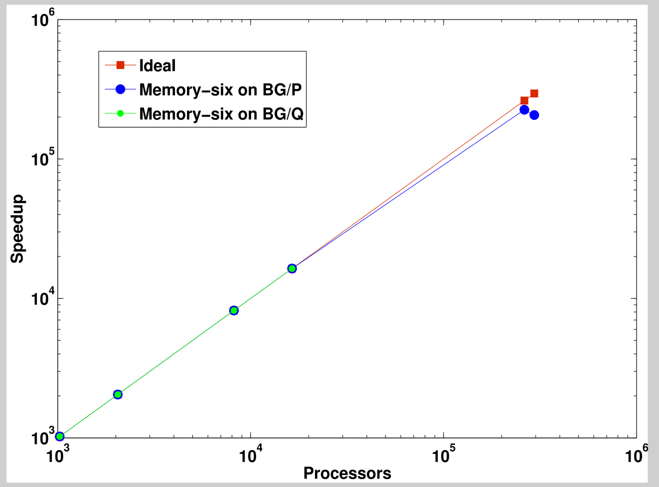

Strong Scaling up to 262,144 processors.

On Blue Gene/Q, we demonstrated equivalent scaling efficiency up to 16,384 processors achieved by running 512 nodes with 32 tasks per node.

Conclusions and Future Work

We have presented a highly-scalable framework for modelling evolutionary game dynamics. Our method enables domain scientists to study population dynamics at an unprecedented scale, spanning a larger population size and greater number memory steps. The excellent weak scaling of our approach enables us to grow the population size dramatically while keeping simulation time reasonable. Our framework will not only allow researchers to assess the role memory plays in game dynamics, but also to determine if there are more complex strategies that lead to the emergence of cooperation between agents. This has the potential to widely broaden the scope of game theory simulation and the fields in which it is used, particularly for large-scale economic models.

References

R. Axelrod and W. Hamilton. The evolution of cooperation. Science, 211(4489):1390, 1981.
M. Nowak and K. Sigmund. Tit for tat in heterogeneous populations. Nature, 355(6357): 250–253, 1992.
M. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. Nature, 364(6432):56–58, 1993.

Acknowledgements: This work was supported by the DOE CSGF, grant number DOE CSGF DE-FG02-97ER25308. This research used resources of Harvard University, Juelich Forschungszentrum, and the ALCF at Argonne National Laboratory, which is supported by the Office of Science of the US DOE under contract DE-AC02-06CH11357. It was funded partially by the DOE Computational Science Graduate Fellowship. Christopher Lee, David Rand, Hanspeter Pfister, Martin Nowak, and Greg Morrisett all contributed input and support.