Leveraging GPUs for Evolutionary Game Theory

Amanda Peters
Harvard University

Outline

- Motivation
- Game Theory 101
- Application Model
- Results

Motivation

- Why does cooperation evolve?

Behavior

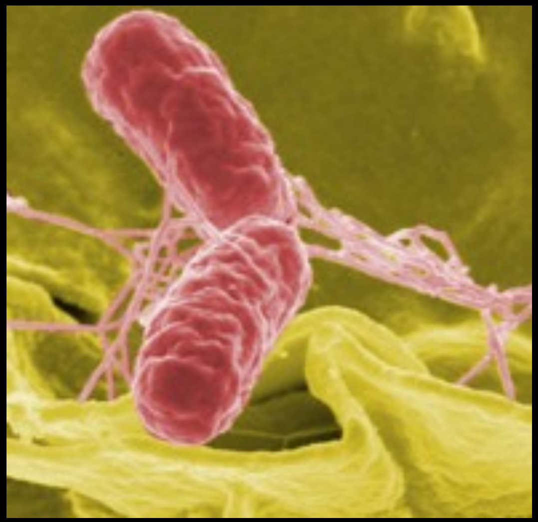

Biology

Economics

- Goal: Create computational model to test role of behavioral strategies and related variables

http://kentsimmons.uwinnipeg.ca/16cm05/1116/16behave.htm
http://gohealed.com/
http://teritoday.thegrocerygame.com/con_Money.cfm?dc=1012&dsc=20235

Contributions

- Heterogenous pipeline
- Enabled the study of behavioral strategies based on many historical steps
- Achieved 209x speedup on 4 Tesla GPU cluster
- Successfully scaled to 72 racks of Blue Gene/P
- Achiever linear scaling up to 262,144 processors

Game Theory 101

Prisoner's Dilemma

Player 2

		C	D
Player 1	C	3,3	0,4
	D	4,0	1,1

Prisoner's Dilemma

		Player 2	
		C	D
Player 1	C	3,3	0,4
	D	4,0	1,1

Rational players will always defect

Tit-for-Tat

- Direct reciprocity
- Unless provoked, the agent will always cooperate
- When provoked, the agent will retaliate
- The agent does not hold grudges
- Iterated PD

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C			
Player 2				

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C			
Player 2	D			

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C	D		
Player 2	$\boxed{D}$			

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C	D		
Player 2	D	D		

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C	D	D	
Player 2	D	$\boxed{D}$		

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C	D	D	
Player 2	D	D	C	

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C	D	D	C
Player 2	D	D	$\boxed{C}$	

Tit-for-Tat Example

	Round 1	Round 2	Round 3	Round 4
Player 1	C	D	D	C
Player 2	D	D	C	C

Win-Stay-Lose-Shift

	C	D
C	R,R	S,T
D	T,S	P,P

- For R&S, TFT and WSLS behave the same
- For T & P, opposite
- TFT promotes adaptation by teaching a lesson, WSLS gets the lesson as well

Problem Size

States

	P1	P2
1	C	C
2	C	D
3	D	C
4	D	D

Problem Size

	States		**Strategies**
	P1	P2	
1	C	C	C or D
2	C	D	C or D
3	D	C	C or D
4	D	D	C or D

Problem Size

	States		Strategies	
	P1	P2		
1	C	C	C or D	C C D
2	C	D	C or D	C D D
3	D	C	C or D	C C D
4	D	D	C or D	C C D

Problem Size

States			**Strategies**			
	P1	P2				
1	C	C	C or D	C	C	D
2	C	D	C or D	C	D	D
3	D	C	C or D	C	C	D
4	D	D	C or D	C	C	D

Problem Size

States			Strategies	
	P1	P2		
1	C	C	C or D	C $\boxed{C}$ D
2	C	D	C or D	C $\boxed{D}$ D
3	D	C	C or D	C $\boxed{C}$ D
4	D	D	C or D	C $\boxed{C}$ D

Problem Size

States			Strategies	
	P1	P2		
1	C	C	C or D	C C $\boxed{D}$
2	C	D	C or D	C D $\boxed{D}$
3	D	C	C or D	C C $\boxed{D}$
4	D	D	C or D	C C $\boxed{D}$

Problem Size

	States		**Strategies**	
	P1	P2		
1	C	C	C or D	C C $\boxed{D}$
2	C	D	C or D	C D $\boxed{D}$
3	D	C	C or D	C C $\boxed{D}$
4	D	D	C or D	C C $\boxed{D}$

Number of States = $2\wedge^{2*\text{memSteps}}$

Problem Size

	States		**Strategies**	
	P1	P2		
1	C	C	C or D	C C D
2	C	D	C or D	C D D
3	D	C	C or D	C C D
4	D	D	C or D	C C D

Number of States = 2^ 2*memSteps

Number of Strategies = 2^ numStates

Problem Size cont.

Memory Steps	Number of Strategies
1	16
2	65536
3	1.84*10^19
4	1.15*10^77

Population Model

- Selection at the individual level
- Replace organism with an individual of differing strategy
- Temperature of selection
- Random mutation
- Zero population growth

Parallel Evolutionary Biology Suite (PEBS)

Why Parallelize it?

- Reduce computational time
- Enable trials of more behavioral strategies
- Enable analysis of different variables:
  - Punishment
  - Kin selection
  - Reciprocity
- **Longer historical recall**

Pipeline

GPU

Large-scale MPI

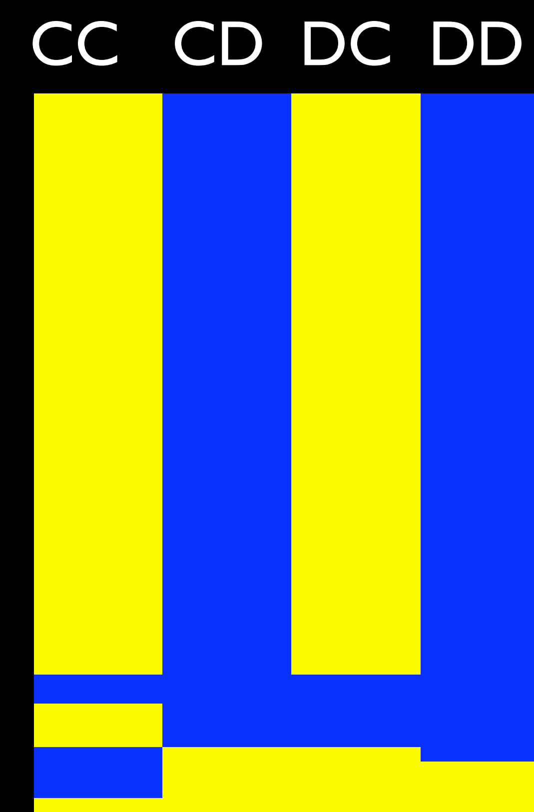

Results

GPU Model

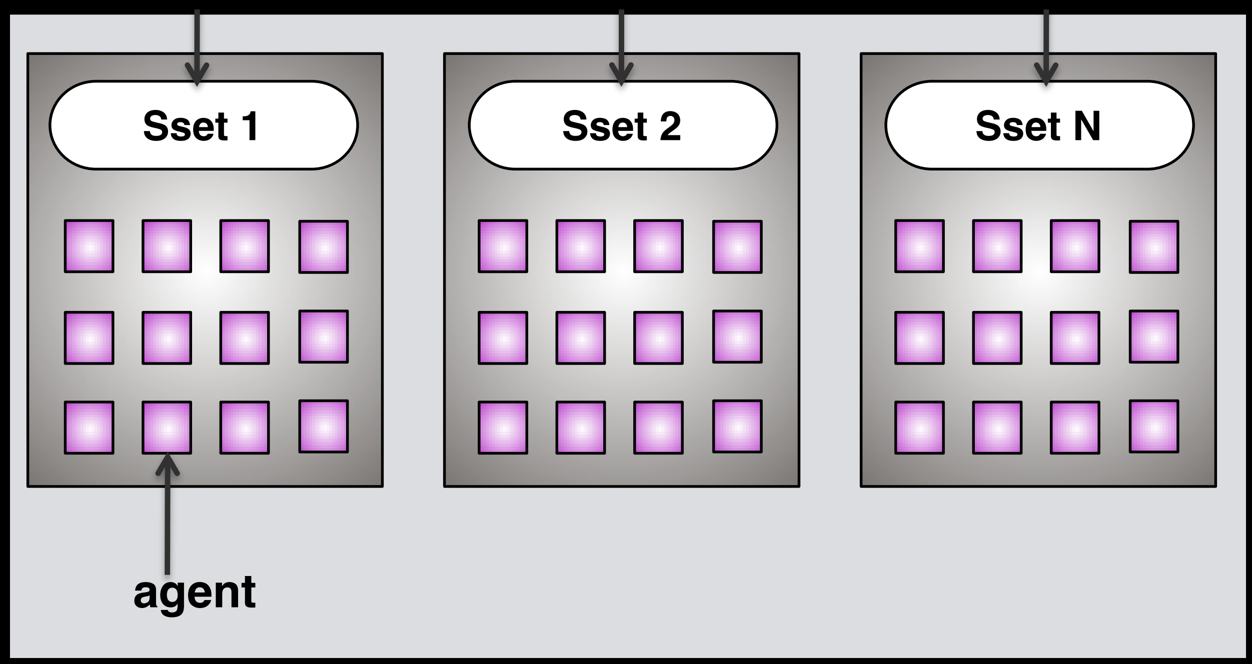

- Each Sset assigned to a block on the GPU
- Each thread handled the interaction between an agent with that strategy and another agent

GPU Process

- Phase 1: GPU kernel call per game

GPU Process

- Phase 1: GPU kernel call per game
- Phase 2: Multiple Block Implementation

GPU Process

- Phase 1: GPU kernel call per game
- Phase 2: Multiple Block Implementation
- Phase 3: All Data processing on the Device

GPU Process

- Phase 1: GPU kernel call per game
- Phase 2: Multiple Block Implementation
- Phase 3: All Data processing on the Device
- Phase 4: Fine Tuning

GPU Process

- Phase 1: GPU kernel call per game
- Phase 2: Multiple Block Implementation
- Phase 3: All Data processing on the Device
- Phase 4: Fine Tuning
- Phase 5: Expand to the GPU cluster

MPI Model

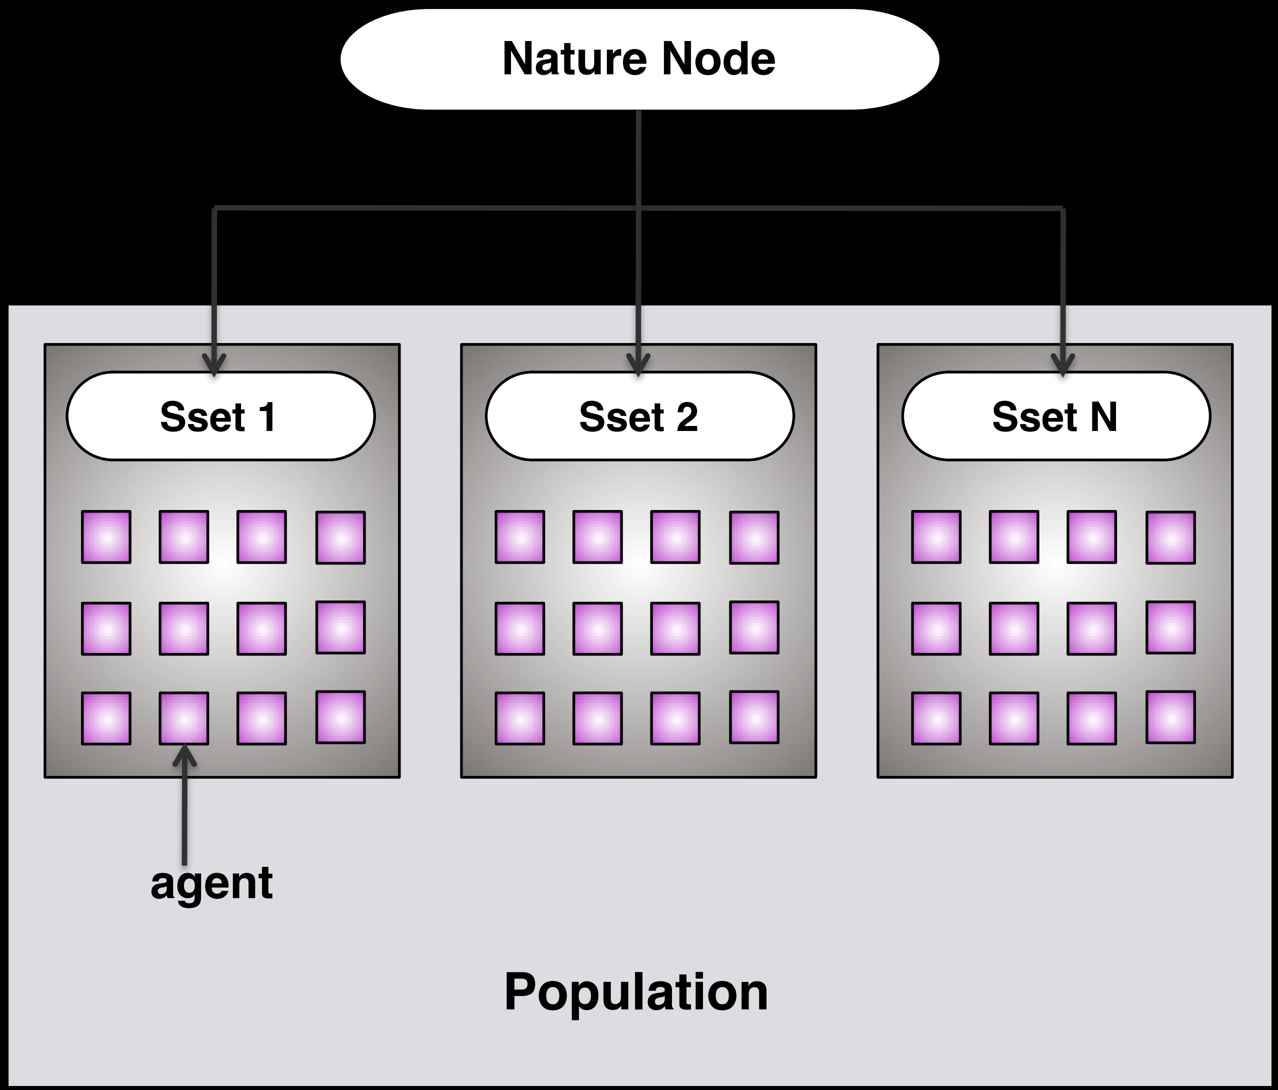

- Large-scale population model
- Mutations, Learning, and Errors
- Heavy communication

MPI Version

- Large-scale population model
- Mutations
- Errors
- Heavy communication

Results

Runtime Minimization

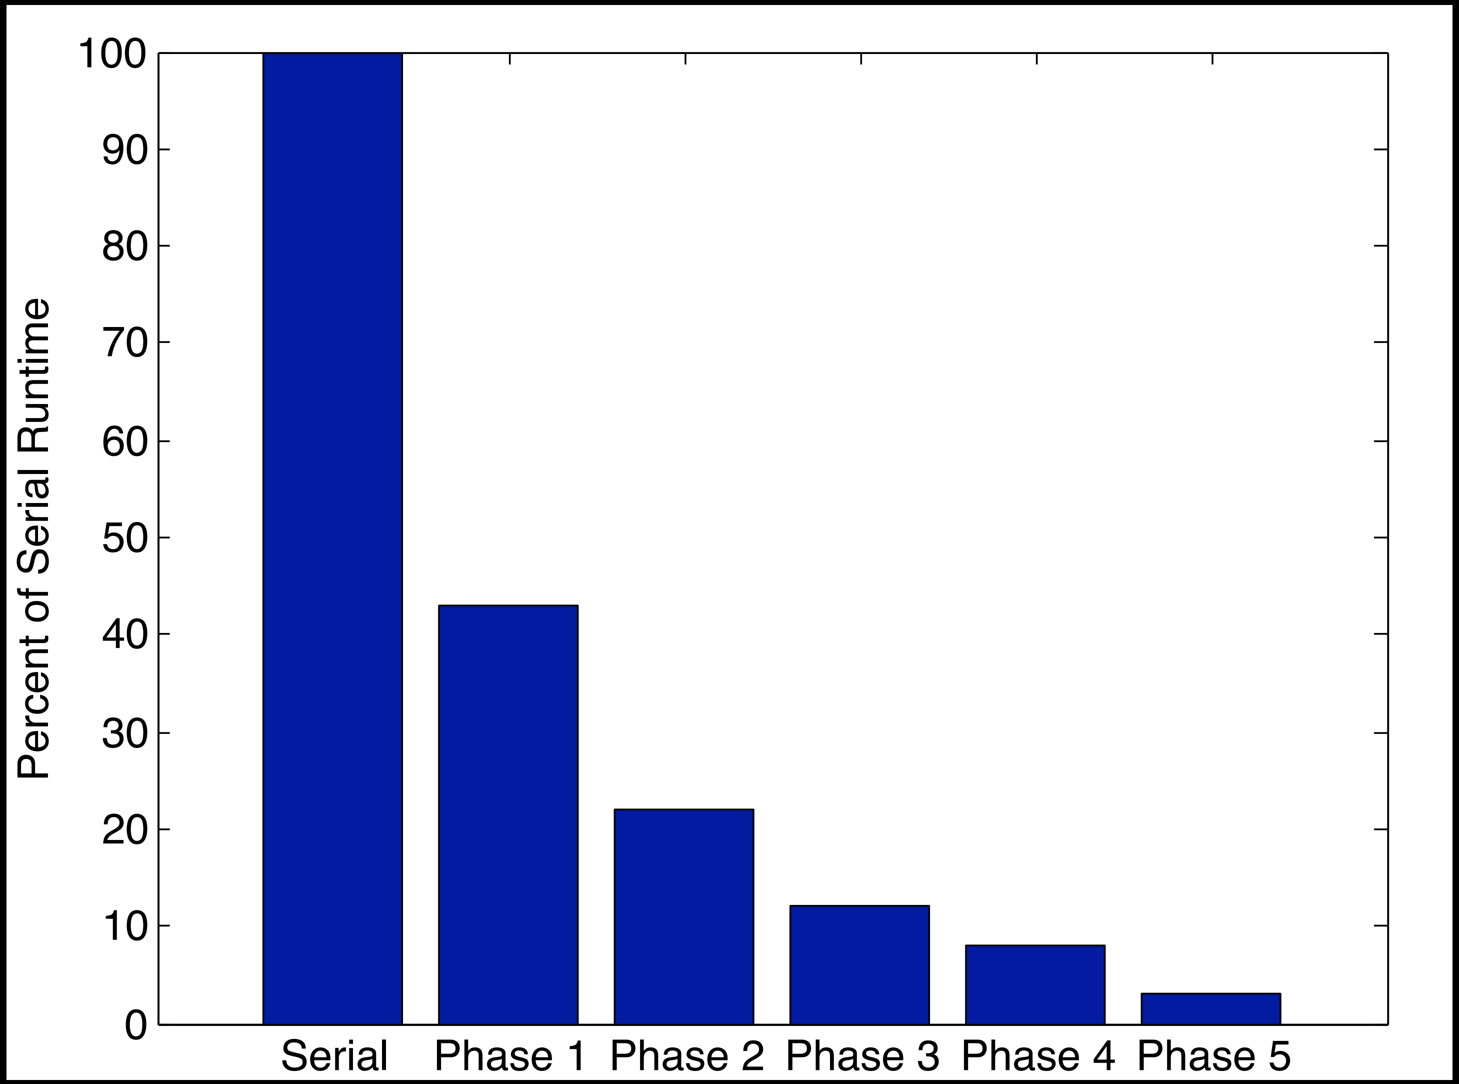

- 97% time reduction

Weak Scaling

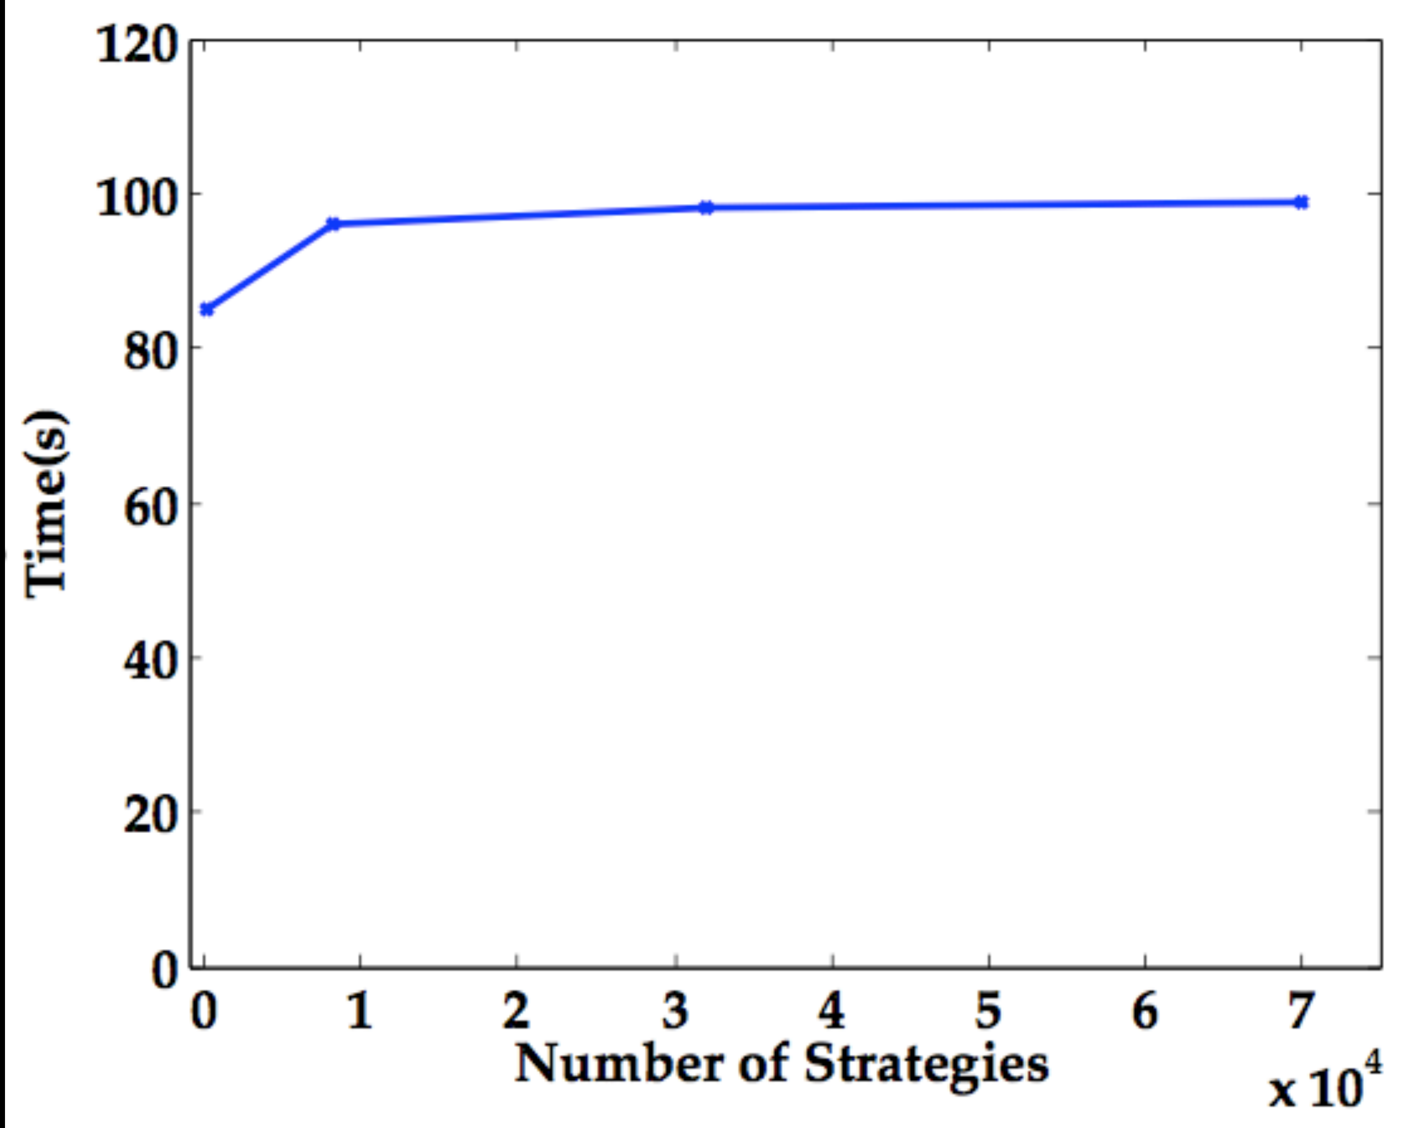

Strong Scaling

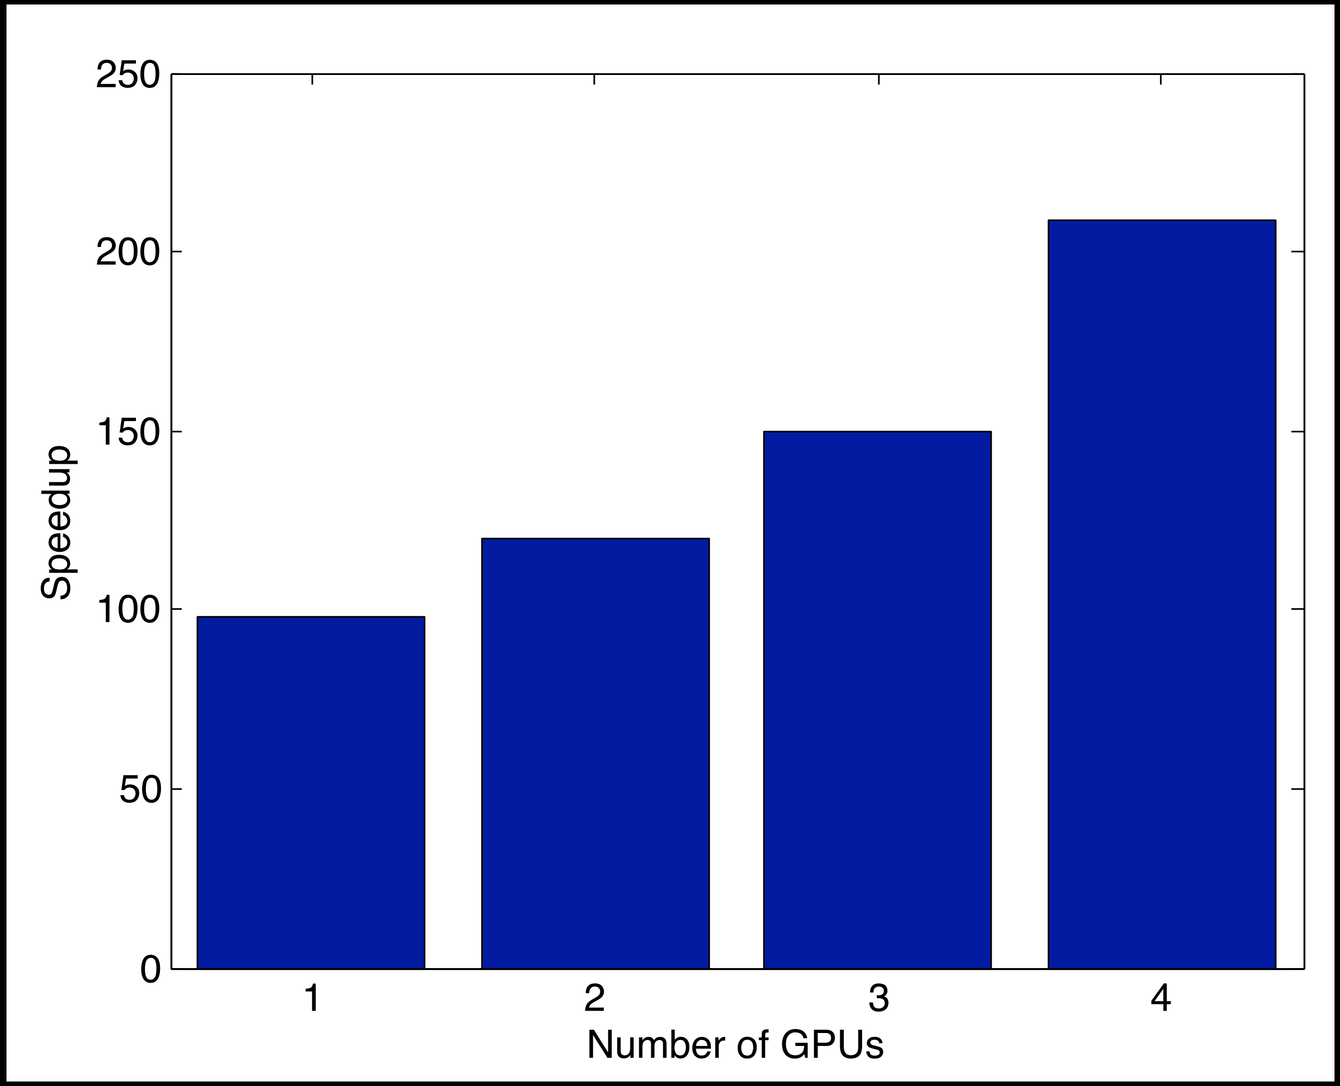

Jugene

- 72 racks of Blue Gene/P
- 294,912 cores
- Peak Performance: 1 Petaflop
- Memory: 2 Gb per core

Blue Gene Results

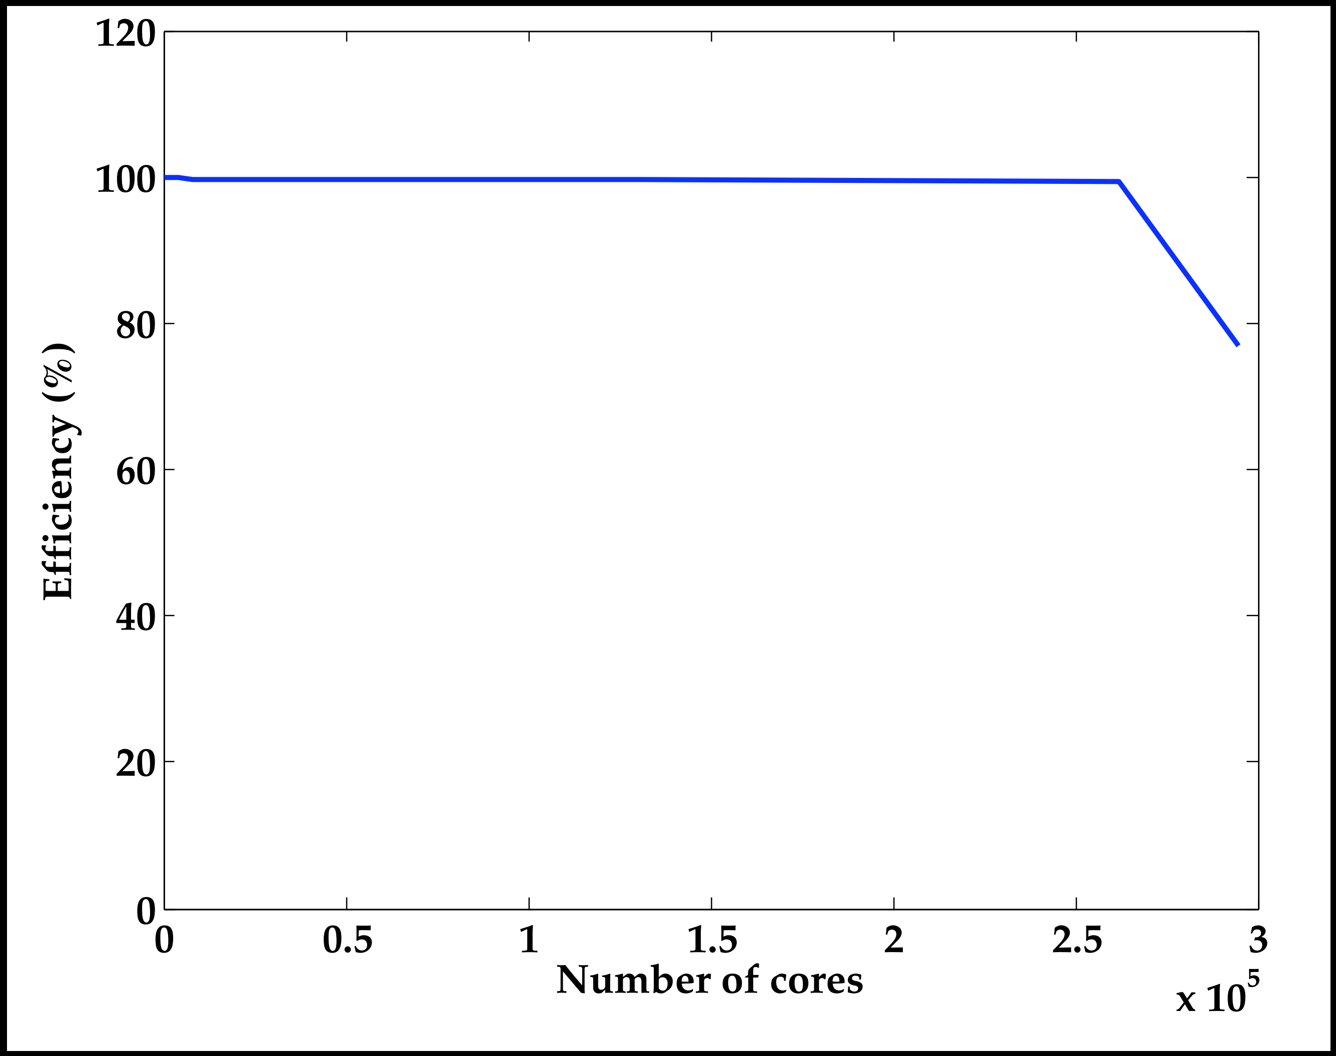

TfT Results

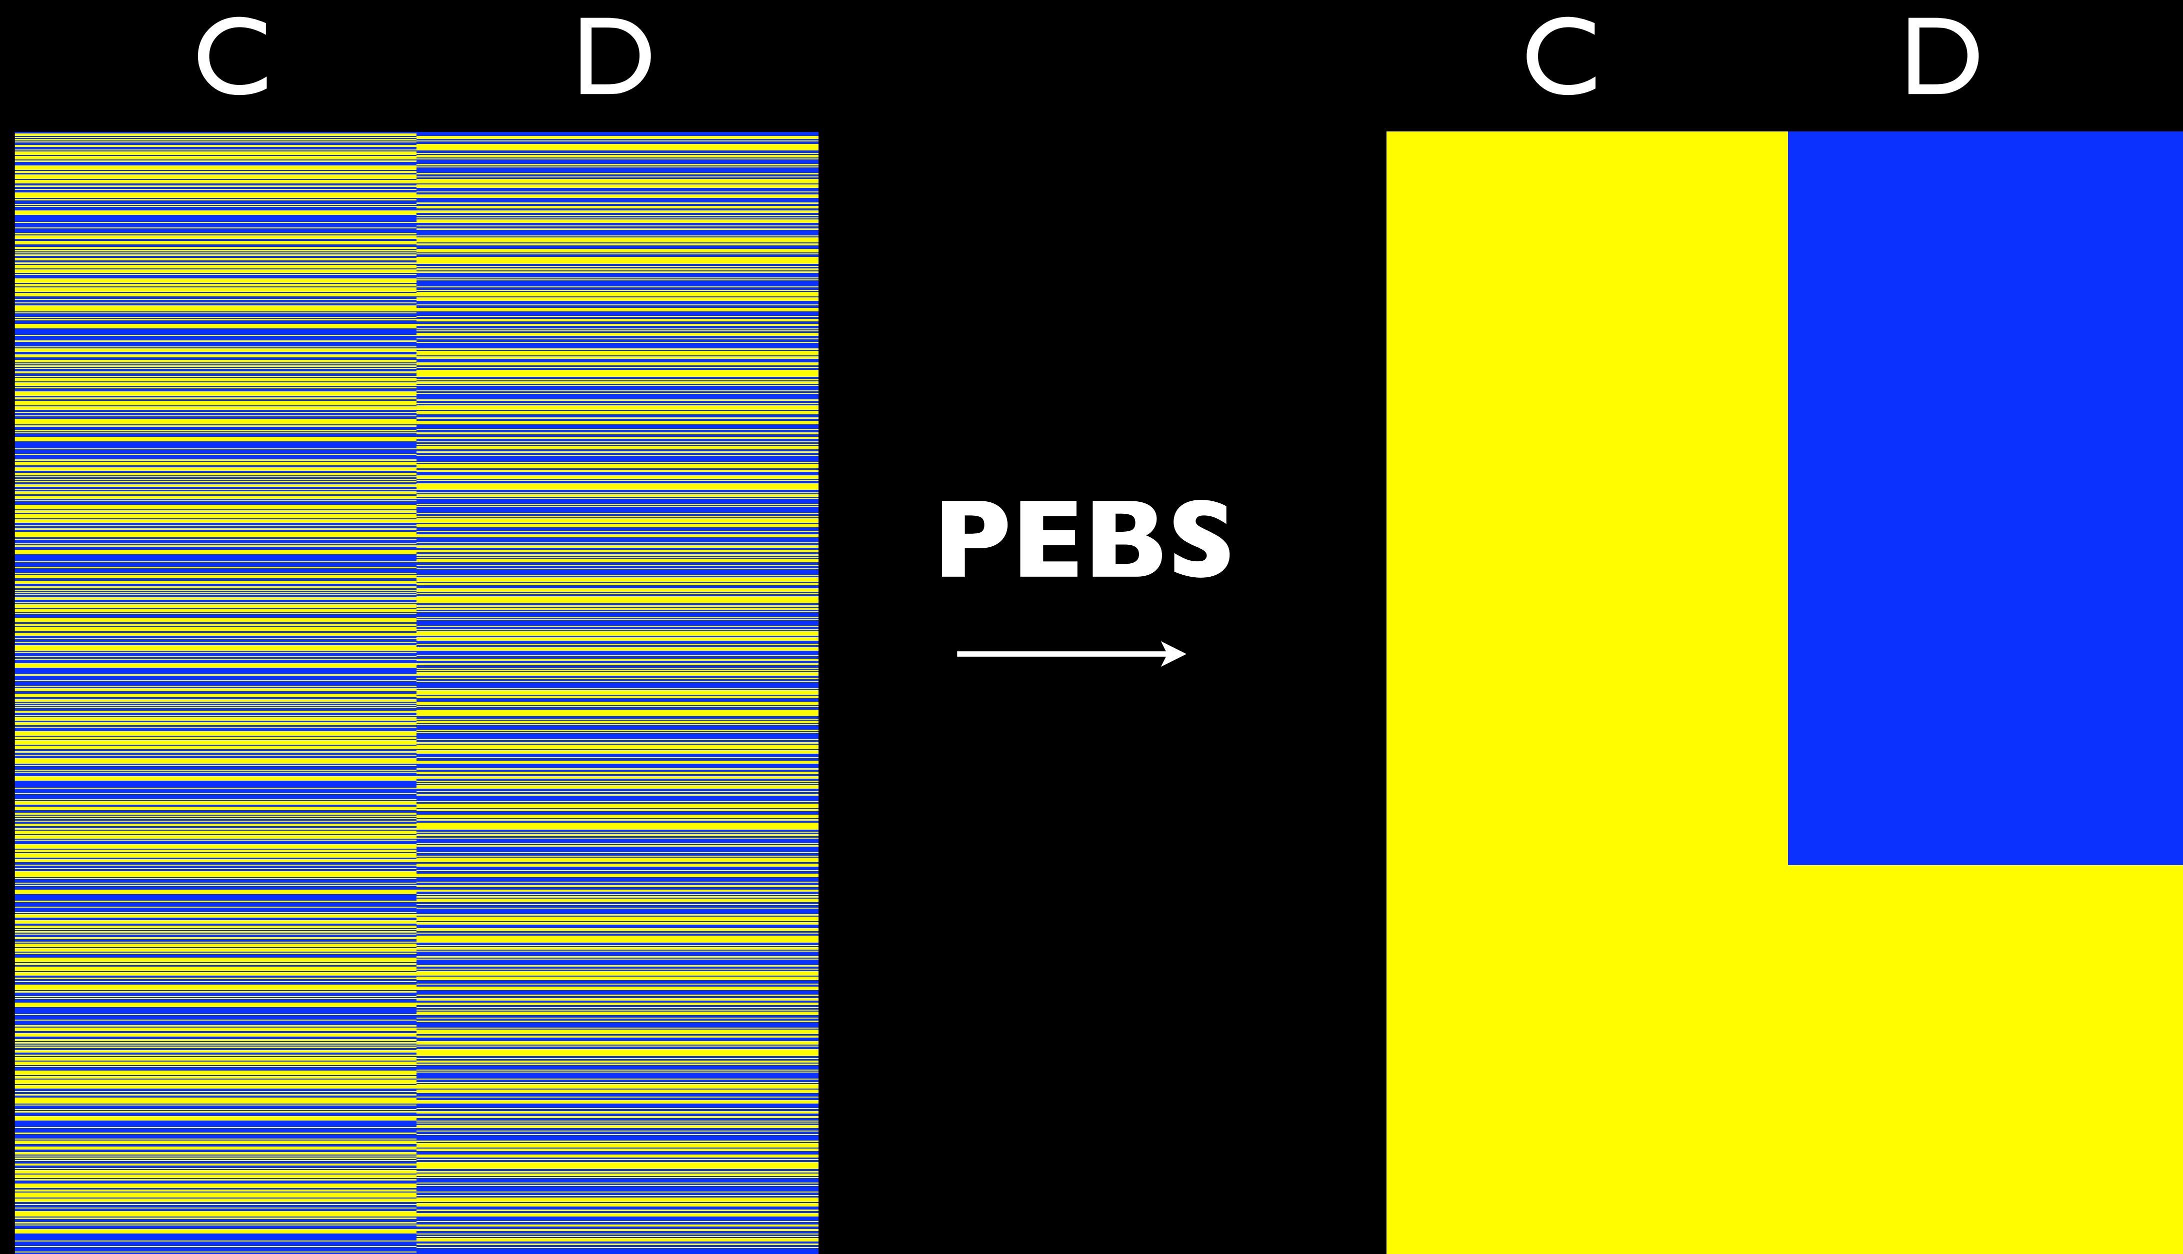

Results WSLS

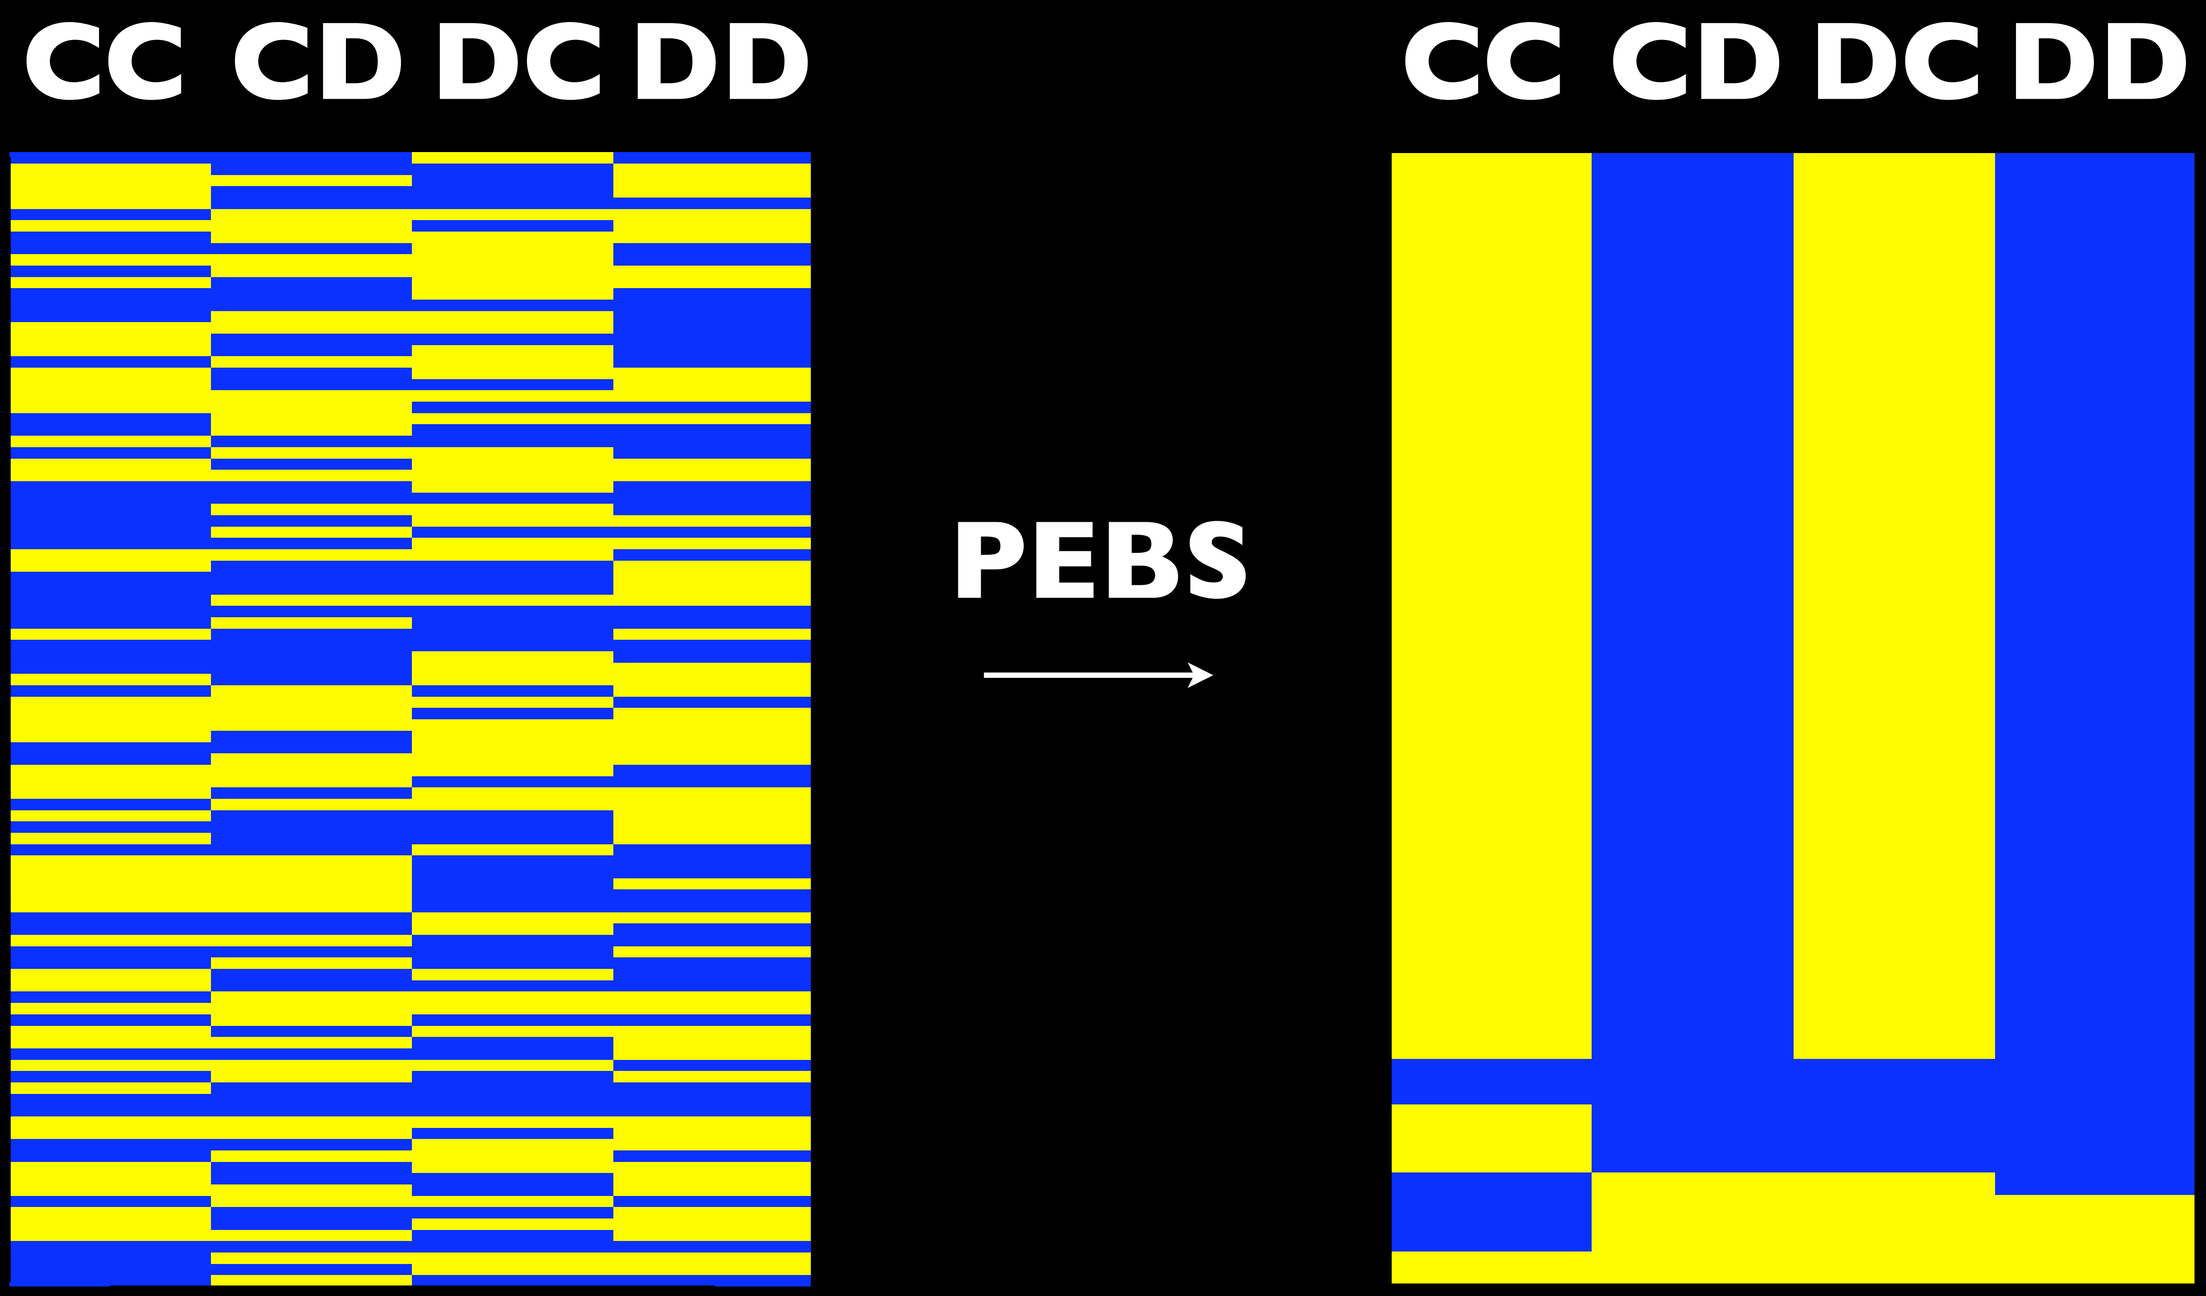

Summary

- Heterogenous data pipeline
- Enabled the study of behavioral strategies based on many historical steps
- Achieved 209x speedup on 4 Tesla GPU cluster
- Successfully scaled to 72 racks of BG/P
- Achiever linear scaling up to 262,144 processors

Future Work

- Move large run to GPU cluster
- Study greater number of memory steps
- Look at biological implications
- Probabilistic strategies

Thank you

- David Rand, Harvard University
- Christopher Lee, Harvard University
- Martin Nowak, Harvard University
- Greg Morrisett, Harvard University
- Hanspeter Pfister, Harvard University
- Joy Sircar, Harvard University

Questions?